

Improving Federal Cybersecurity Governance Through Data-Driven Decision Making and Execution

Douglas Gray
Julia Allen
Constantine Cois
Anne Connell
Erik Ebel

William Gulley
Michael Riley
Robert Stoddard
Marie Vaughan
Brian D. Wisniewski

September 2015

TECHNICAL REPORT
CMU/SEI-2015-TR-011

CERT Division

Distribution Statement A: Approved for Public Release; Distribution is Unlimited

<http://www.sei.cmu.edu>



Copyright 2015 Carnegie Mellon University

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Department of Homeland Security or the United States Department of Defense.

This report was prepared for the

SEI Administrative Agent
AFLCMC/PZM
20 Schilling Circle, Bldg 1305, 3rd floor
Hanscom AFB, MA 01731-2125

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Carnegie Mellon® and CERT® are registered marks of Carnegie Mellon University.

DM-0002528

Table of Contents

Abstract	v
1 Introduction	1
1.1 Background	1
1.2 Objective	2
1.3 Target Audience	2
2 The Observe, Orient, Decide, Act (OODA) Loop	5
2.1 What Is the OODA Loop?	5
2.2 How the OODA Loop Compares to Other Cycle Approaches	6
2.3 Why Use the OODA Loop?	7
3 What is Cybersecurity Governance?	9
3.1 How Is Governance Different from Operations?	9
3.2 Facets of Cybersecurity Governance	11
3.2.1 Doctrine and Strategy	12
3.2.2 Enterprise Portfolio Management	13
3.2.3 Financial Resource Management	13
3.2.4 Enterprise Acquisition and Materiel Management	13
3.2.5 Human Resources Management and Leader Development	13
3.2.6 Organizational Structure Management	14
3.2.7 Organizational Training and Awareness	14
3.2.8 Legal, Regulations, Investigations, and Compliance	14
3.2.9 Enterprise Risk Management	15
4 Enabling Data-Driven Decision Making	16
4.1 Collecting Situational Awareness Data and Information (Observe)	16
4.1.1 Authoritative vs. Non-Authoritative Data Sources	17
4.1.2 Data Engineering	18
4.1.3 Management Information	18
4.1.4 Threat Information	19
4.1.5 Automated Vulnerability Information	20
4.2 Position the Enterprise for Action (Orient)	21
4.2.1 Defining Environmental Constraints for Implementation of the OODA Loop	22
4.2.2 Data Science	24
4.2.3 Visualization	24
4.2.4 Prioritization of Problems to be Addressed	26
4.3 Key Planning and Decision-Making Factors (Decide)	28
4.3.1 Determine Time Available for Decision Analysis	29
4.3.2 Determine Hypothesis or Theory	29
4.3.3 Determine Enablers	30
4.3.4 Determine Criteria and Weighting	30
4.3.5 Determine Courses of Action	31
4.3.6 Evaluate Courses of Action	31
4.4 Enabling Success at the Point of Execution (Act)	32
5 Implementation Using a Maturity Model	35
6 Conclusion	36
Appendix A: Collecting and Categorizing Threat Information	37
Appendix B: Automated Vulnerability Collection	41

Appendix C: Use of Indices and Quantitative Methods to Enhance Cybersecurity Governance Decision-Making Insights and Lessons from Data Science Literature	44
Appendix D: Mapping of Facets of Cybersecurity Governance to Other Frameworks	59
Bibliography	60

List of Figures

Figure 1:	Target Audience for this Technical Report	4
Figure 2:	The OODA “Loop” Sketch [Boyd 1996]	5
Figure 3:	Comparison of OODA Loop to Other Cycle Processes	6
Figure 4:	The OODA “Loop” Sketch [Boyd 1996]	8
Figure 5:	Cybersecurity Governance and NIST CSF Notional Organizational Information and Decision Flows [NIST 2014, pg. 12]	10
Figure 6:	Facets of Cybersecurity Governance	12
Figure 7:	Matching Prevailing Attack Patterns to an Organization’s Target Surface Area	20
Figure 8:	Sources of Mandates and Constraints	23
Figure 9:	Problem Prioritization Decision Tree	27
Figure 10:	Transition from Decide to Act Based on a Hypothesis	33
Figure 11:	Transition from Decide to Act Based on a Theory	34
Figure 12:	Weighted Product Model	55

List of Tables

Table 1:	Security Controls Grouped into OODA Categories	6
Table 2:	Comparison of Operations and Governance	11
Table 3:	Usability Heuristics	26
Table 4:	Decision Matrix	32
Table 5:	Mapping of Components, Attributes, and Desired Threat Capabilities	37

Abstract

Although efforts are underway through Information Security Continuous Monitoring initiatives to improve situational awareness and risk mitigation at the operational level, the federal government must make better enterprise-level cybersecurity decisions in the shortest time possible. This report outlines an approach called Data Driven Cybersecurity Governance Decision Making. This approach leverages the Observe, Orient, Decide, Act (OODA) loop used by the U.S. Department of Defense to enable decision makers at the strategic levels of government to best set the conditions for success at the point of execution. To best target the unique considerations of enterprise decision makers, this report discusses the difference between cybersecurity governance and cybersecurity operations. Within this context, it describes best practices in collecting and analyzing authoritative data present in the federal space to develop a level of situational awareness tailored to decision makers' needs in a cybersecurity governance scorecard. Cybersecurity governance decision makers can leverage this enhanced situational awareness to support a data-driven decision-making process that targets root causes of the problems facing the federal government enterprise. Finally, the report discusses key considerations to ensure success at the point of execution based on work performed in the Observe, Orient, and Decide phases of the OODA Loop.

1 Introduction

1.1 Background

With six million employees and service members and over 126 departments and agencies (D/A) the United States Government stands at a distinct disadvantage with respect to defending against cybersecurity threats.¹ Quintessentially, threats hold the initiative and force federal cybersecurity defenders to react within the confines of their capabilities. With the constantly increasing rate of change in information-technology (IT) and practices, system defenders can naturally be expected to lose ground in their defenses. In addition, the federal government is further constrained by decision-making and execution processes that are rooted in multi-level statutory, judicial, and executive decision making that seek to ensure accountability to the democratic process but operate at substantially less than machine speed. What's more, the federal government is fragmented, with approximately 70 percent of all government agencies consisting of small organizations averaging 150 employees (FedScope), stretching the government's ability to defend all of its IT functions in such a heterogeneous environment.

The U.S. Government faces increasingly complex and constantly changing business and operational environments. Federal government agencies are constantly bombarded with conditions and events that can introduce stress and uncertainty that may disrupt their effective operation and preclude them from achieving their missions.

Stress to the federal government's ability to operate securely and continue to achieve its mission in the presence of cybersecurity incidents can come from many sources. For example

- Pervasive use of technology and technology advances (such as cloud computing and the proliferation of mobile devices) are helping agencies automate business processes and make them more effective in achieving their missions. However, the cost to agencies is that the technology often introduces operational complexity, takes specialized support and resources, and creates an environment that is rife with security vulnerabilities and risks.
- Agencies increasingly depend on public/private partnerships and service providers to achieve their missions. External partners provide essential skills and functions, with the aim of increasing productivity and reducing costs. As a result, an agency exposes itself to new risk environments inherent in partner organizations, their risks, and their supply chains. By employing multiple partners to provide a government service, an agency cedes control of mission assurance in exchange for cost savings.
- The increasing globalization of agency missions and the companion supply chains pose a problem in that governance and oversight must cross governmental, organizational, geographical, and jurisdictional lines as never before. In addition, the emerging worldwide sociopolitical environment is forcing agencies to consider threats and risks that have previously not been on their radar screens. Recent well-publicized events have changed the view of what is feasible and has expanded the range of outcomes that an agency must attempt to prevent and from which it must be prepared to recover.

¹ These employees and service members are part of FedScope, the Defense Manpower Data Center, and HR.

When one couples this changing environment with the sheer immensity and complexity of government, a seemingly insurmountable problem set emerges. However, even with these disadvantages, there are opportunities. The size and resourcing of the federal government (in 2013, the federal government spent more on IT than the 2012 gross domestic product of Latvia²) opens opportunities for information sharing and means the government can allocate resourcing for well-targeted expenditures. However, what is required is a methodology at the macro level to gain situational awareness, make decisions, and effectively execute those decisions. Without such an ability, federal departments and agencies (D/As) will continue to lose ground to threat actors in the cybersecurity domain.

In order to turn the federal government's immense size into an advantage, the problem and solution sets must be decomposed. Today the body of available cybersecurity knowledge is saturated with best practices on managing risks at a direct, operational level. However, best practices for managing cybersecurity at the indirect, strategic level are few and far between. There is little information on what capabilities are available to decision makers at the enterprise level and how they can be brought to bear to set the context so that operational cybersecurity can be successful.

1.2 Objective

This technical report focuses on cybersecurity at the indirect, strategic level. In order for cybersecurity decision makers at the tactical or implementation level to be successful, a supportive contextual environment must be established. This report seeks to deconstruct cybersecurity governance and provide a framework within which the federal government can

- obtain data relevant to effective cybersecurity governance
- transform that data into information that enables cybersecurity governance decision makers to make sense of their environments
- enable effective and efficient decision making
- enable success at the point of execution of those decisions
- do all of the above in a way that is faster and more effective

To do this, we will use the Observe, Orient, Decide, Act (OODA) Loop as the basis. This report will show how the OODA-based approach can be used as a means to capture status and prioritize actionable information in a real-time cybersecurity governance scorecard that leverages authoritative and non-authoritative data sources. Additionally, we will demonstrate how OODA-based, data-driven cybersecurity governance can be consistent with other frameworks such as the NIST Risk Management Process (RMP) and the Deming Management Method's Shewhart Cycle, with the potential for use in providing metric input to progression and capability maturity models.

1.3 Target Audience

This technical report is intended for use by cybersecurity governance practitioners and decision makers with portfolio and program-management responsibilities for cybersecurity in the federal government space. These decision makers generally have responsibility to set the environment

² Federal IT Dashboard (<https://www.itdashboard.gov/sites/default/files/exhibit53report/4>), World Bank GDP (http://data.worldbank.org/indicator/NY.GDP.MKTP.CD?order=w_bapi_data_value_2012+w_bapi_data_value+w_bapi_data_value-last&sort=desc)

within which cybersecurity operations can proceed effectively. Examples of such audience members include

- members of Congress
- Congressional personnel and committee staff members
- members of the Office of Management and Budget (OMB) Office of Information & Regulatory Affairs
- members of OMB Office of E-Government & Information Technology
- members of the National Security Council Staff
- staff of the Program Manager for the Information Sharing Environment
- staff of the Department of Homeland Security with cybersecurity oversight including, but not limited to
 - the National Cybersecurity and Communications Integration Center (NCCIC)
 - the Federal Network Resilience Division
- staff of the National Institute of Standards and Technology (NIST)
- staff of the Office of the Secretary of Defense with cybersecurity oversight
- staff of United States Cyber Command and cyber-related service component commands and field-operating agencies
- staff of the National Security Agency's Information Assurance Directorate
- staff of the Office of the Director of National Intelligence with cybersecurity oversight
- staff of the Government Accountability Office
- D/A chief information officers and chief information security officers and their staffs
- D/A inspectors general and their staff
- staff of the General Services Agency with oversight of cybersecurity-related programs
- D/A authorizing officials, information assurance program managers, and information assurance managers
- members of the private sector supporting the audience members listed above

Figure 1 below graphically depicts the audience for this technical report.

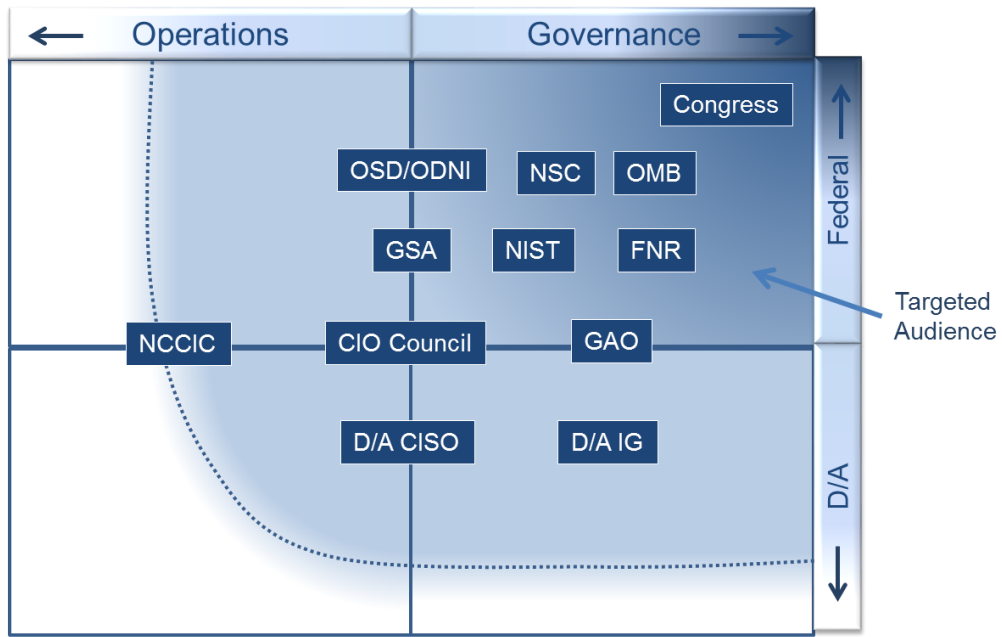


Figure 1: Target Audience for this Technical Report

2 The Observe, Orient, Decide, Act (OODA) Loop

2.1 What Is the OODA Loop?

The OODA Loop is a mental model for conceptualizing how individuals and organizations make decisions. It was originally developed by Colonel John Boyd, a U.S. Air Force fighter pilot, based upon his experiences as a veteran of the Korean War and through his study of strategic theorists, such as Sun Tzu, Julian Corbet, T. H. Lawrence, J. F. C. Fuller and Basil Liddell Hart [Osinga, 2007].

The OODA Loop was Boyd's attempt to explain how we create mental patterns or "concepts of meaning" in order to be able to comprehend and cope with our environment. Boyd described this approach as an attempt "to sketch out how we destroy and create these patterns to permit us to both shape and be shaped by a changing environment. In this sense, the discussion also literally shows why we cannot avoid this kind of activity if we intend to survive on our own terms" [Boyd 1976]. Throughout the years since its creation, Boyd's work has been used within the Department of Defense, particularly within the United States Marine Corps, as the foundation for much of the doctrine of maneuver warfare and within the broader business, legal, and analytics communities. Greg Wood's blog post from May 2012 on Solving Business Problems with Dogfights and OODA Loops highlights its specific use within High Performance Analytics [Wood 2012].

Figure 2 shows the OODA Loop Flow Diagram as Boyd depicted in his OODA "Loop" Sketch [Boyd 1996].

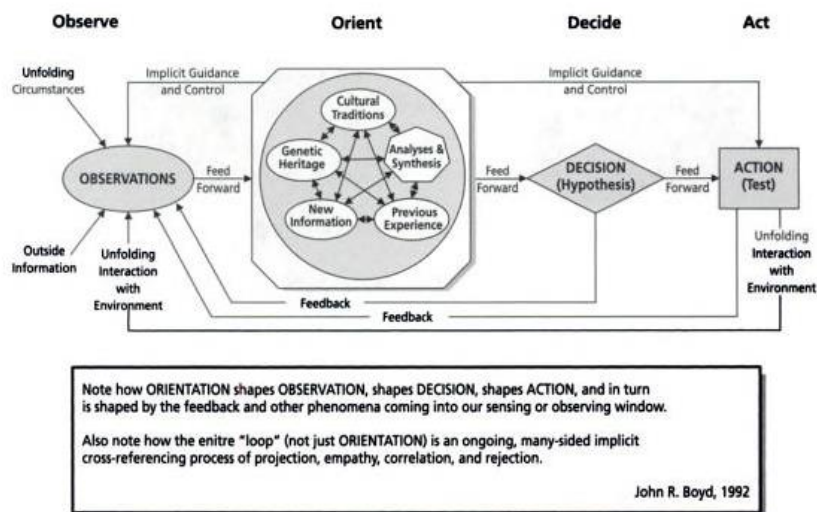
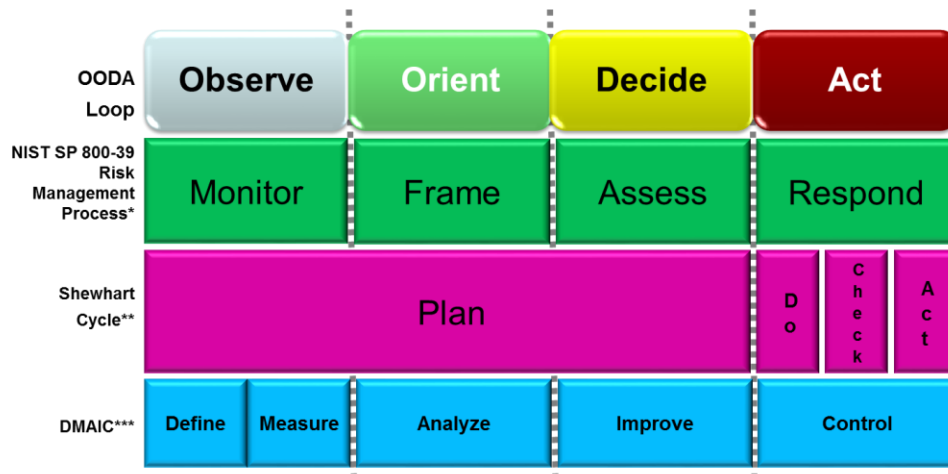


Figure 2: The OODA "Loop" Sketch [Boyd 1996]

The OODA Loop, in the military context, describes the ability to acquire, process and act upon information in comparison to one's adversary's ability to do so. The common phrase, "getting inside their decision cycle," is a reference to being able to cycle through this loop faster than your adversary.

2.2 How the OODA Loop Compares to Other Cycle Approaches

The OODA Loop provides a fairly robust framework for addressing cybersecurity within an organization. The graphic below outlines the comparison of Boyd's OODA Loop and other common cycle approaches.



*Source: NIST SP 800-39. According to NIST SP 800-39, the Risk-Management Process is not a sequential process like the OODA Loop or the Shewhart Cycle. All components can receive input and send output directly to all other components.

**Source: Walton (1988)

***Source: Lean Six Sigma Corporation

Figure 3: Comparison of OODA Loop to Other Cycle Processes

The NIST RMP provides a key example of how Boyd's ideas can be leveraged within a new domain. The NIST Special Publication (SP) 800-53 Security Controls can be grouped into OODA categories in order to allow planners to leverage the existing framework but look at it in a new light [Hale 2012].

Table 1: Security Controls Grouped into OODA Categories

OODA Loop Category	NIST SP 800-53 Security Controls	Example Practices
Observe	CA-7 Continuous Monitoring	Continuous monitoring program that includes configuration management process; determination of security impact of changes; ongoing security control assessments; reporting security state to appropriate officials
	SI-3 Malicious Code Protection	Employ, update, and configure malicious code protection.
	SI-4 Information System Monitoring	Deploy monitoring devices to monitor events, detect attacks, and identify unauthorized use.
	SI-7 Software and Information Integrity	Detect unauthorized changes to software and information.
	AU-6 Audit Review, Analysis, and Reporting	Review and analyze audit records for indication of inappropriate or unusual activity.
	RA-5 Vulnerability Scanning	Scan for vulnerability in system and applications.
Orient	IR-4 (CE-4) Incident Handling	Correlate incidents with responses (NIST SP800-61: Computer Security Incident Handling Guide).
	RA-3 Risk Assessment	Assess risk, document and review results, and update as changes occur.

OODA Loop Category	NIST SP 800-53 Security Controls	Example Practices
Decide	CM-4 Security Impact Analysis	What are the impacts of changes to the system?
Act	IR-4 (CE-2) Incident Handling	Dynamically reconfigure the system as a result of incidents.
	IR-4 (CE-3) Incident Handling	Handling identify classes of incidents and take appropriate actions as a result of incident class
	SI-4 (CE-3) Information System Monitoring	Employ automated tools to integrate intrusion detection tools into access control and flow control for rapid response.
	SI-2 Flaw Remediation	Identify, report, and correct system flaws.

A key difference between the OODA Loop and the NIST RMP and Shewhart Cycles lies in the numerous discrete feedback and feed forward connections that are identified within the full diagram Boyd developed. This underlying matrix of connections makes the OODA Loop an effective model for understanding complex challenges such as cybersecurity.

2.3 Why Use the OODA Loop?

The OODA Loop is uniquely suited to cybersecurity. It acknowledges the underlying complexity and relationships between the core Observe, Orient, Decide, and Act steps similar to the competing priorities that exist between business (or mission) owners, system owners, network administrators, network defenders, policy and audit analysts, and external partners or service providers. Organizations faced with defending such a complex environment must prioritize their efforts. The OODA Loop allows decision makers to focus on those steps necessary to “improve our ability to shape and adapt to unfolding circumstances, so that we (as individuals or as groups or as a culture or as a nation-state [or in the case of cybersecurity – as a system or enclave]) can survive on our own terms” [Boyd 2006].

Applying the OODA Loop to cybersecurity governance allows an organization to respond in a more agile manner to emerging challenges or threats. These can include traditional cybersecurity threats such as a person or software or environmental occurrence, but in the context of governance at the macro level, may also involve changes in policy, funding or staffing along with increased reporting and compliance requirements. An organization can view these challenges through the lens of Observe, Orient, Decide, and Act as a way to understand the key external drivers and changing circumstances that may necessitate such shifts. As highlighted in Figure 4, these types of constantly evolving issues are incorporated into Boyd’s model through the integrated feedback and feed forward mechanisms.

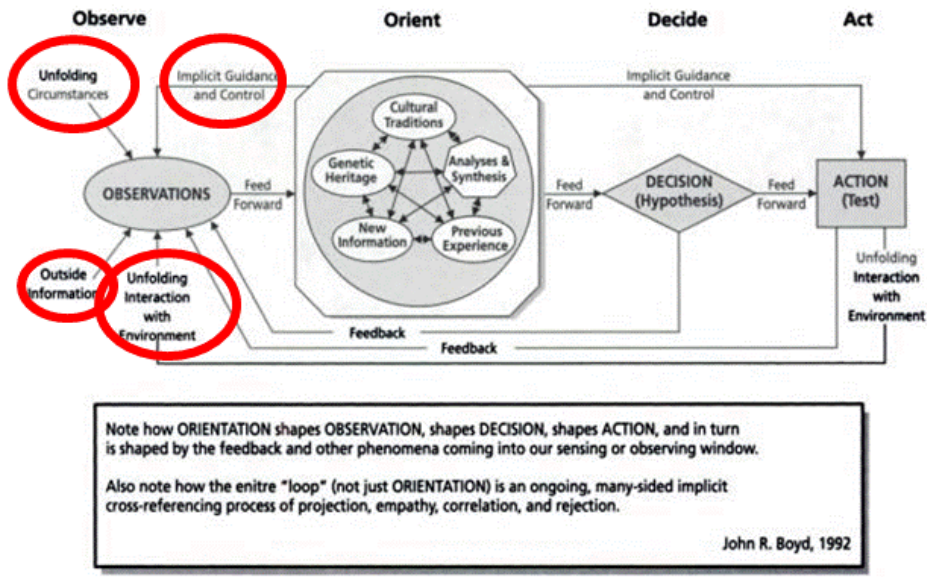


Figure 4: The OODA "Loop" Sketch [Boyd 1996]

A mature organization can leverage this process and easily integrate it into the existing culture. The tolerance for risk within the leadership, whether explicitly defined through a formal process or inferred through informal guidance, can provide a threshold as an organization integrates the variety of information and circumstances within the Observe step and considers it in the context of the Orient step. All of this information, including internal and external indicators, implicit and explicit guidance within the organization, and the context of the circumstances under consideration, leads to the point at which a decision is made. The threshold for this decision will be based upon the collective inputs, guidance, controls, policies, and regulatory requirements. Once this threshold has been reached, an organization is poised to Act. As outlined in Boyd's diagram, one can consider this process of deciding and acting as defining a hypothesis and then testing it and evaluating the outcome of that test.

The OODA Loop provides an abstract yet easily understood approach that can be applied to cybersecurity and integrated into a broader governance process within an organization.

3 What is Cybersecurity Governance?

Because cybersecurity is a federal government-wide concern, it must be governed. Governance includes setting clear expectations for the conduct of the federal government, and directing, controlling, and strongly influencing government D/As to achieve these expectations. It includes specifying a framework for observing, orienting, deciding, and acting, with assigned decision rights and accountabilities, intended to consistently produce desired behaviors and actions. Governance relies on well-informed decision making and the assurance that resulting decisions are routinely enacted as intended. Governance is most effective when it is systemic and woven into the culture and fabric of federal-government behaviors and actions.

Governance and risk management are inextricably linked—governance is an expression of responsible risk management and effective risk management requires efficient governance. Inserting cybersecurity into ongoing governance and risk management conversations is an effective and sustainable approach for addressing security. To achieve a sustainable capability, the federal government must make governing cybersecurity the responsibility of senior leadership, not of other roles that lack the authority, accountability, and resources to act and enforce compliance.

Thus the focus and direction for governing cybersecurity must come from the highest levels within the federal government. The program for achieving an acceptable level of security must be adequately promoted (fostering a security-aware culture), resourced, and monitored and managed in the same fashion as any other mission-critical program or service. Active and visible leadership, sponsorship, and oversight are necessary to ensure that the federal government is achieving its goals as expected. Governing security must be aligned with and support the achievement of the government's strategic objectives. Focusing on these objectives provides the rationale for investing in cybersecurity activities—because they enable the federal government to achieve its mission. Without each of these actions, the government and its agencies will not achieve its desired security posture and will likely fall short in its ability to adapt to, respond to, and recover from disruption and stress and thus to continue to provide mission-critical services during normal and disrupted operations.

3.1 How Is Governance Different from Operations?

Simply stated, governance is oversight—as contrasted with operations, which focuses on implementation and execution. Some of the fundamental activities that governance and operations perform, and the distinctions between them, include the following:

- Governance establishes doctrine, strategies, objectives, key performance indicators, and the policies, plans, and programs to achieve these. These are further refined, decomposed, and allocated to agencies and their operating units, which are responsible for their implementation, tracking, and performance measurement.
- Governance commits funding and resources. Operations implements plans and programs based on these resources and provides feedback on their sufficiency to achieve plans and programs.

- Governance identifies critical operational risks and strategies for effectively mitigating them. Operations implements mitigation actions and controls, provides feedback on how well they are working, and identifies new risks as well.
- Governance establishes compliance requirements. Operations works to meet these requirements, providing feedback when this is not possible within the allocated funding and resources.
- Governance promotes a security-aware culture. Operations personnel conduct themselves in a manner consistent with this culture.

In its Cybersecurity Framework (CSF), NIST uses three levels to describe a common flow of information and decisions within an organization. In this context, cybersecurity governance would refer to actions taken at the executive and business/process levels (see Figure 5).

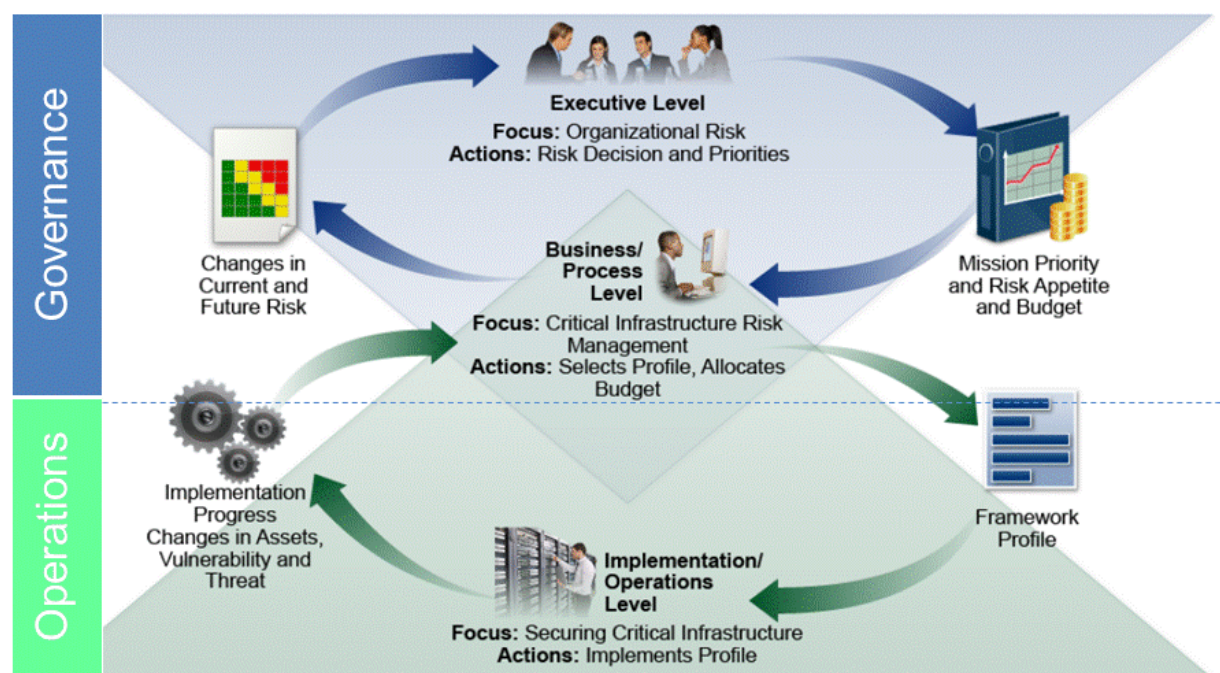


Figure 5: Cybersecurity Governance and NIST CSF Notional Organizational Information and Decision Flows [NIST 2014, pg. 12]

Additional distinctions between governance and operations can also be drawn with respect to scope of concern, timescale of decisions and actions, level of abstraction for activities, and management impact as shown in Table 2.

Table 2: Comparison of Operations and Governance

	Operations	Governance
Scope	Individual networks, systems, users, agencies, organizations	Multiple networks, systems, user bases, agencies, organizations
Timescale	Immediate to 6 months	6 to 36 months ³
Level of Abstraction	Transactional	Trends, aggregations
Management Impact	Direct interaction	Context setting

The next section further elaborates the scope of cybersecurity governance by describing nine facets that are typically within the scope of such governance action.

3.2 Facets⁴ of Cybersecurity Governance

The Facets of Cybersecurity Governance are definable aspects that, in the aggregate, represent a reasonable set of areas of concern for governance decision and action [U.S. Navy 2012]. These facets are shown in Figure 6, in relation to the OODA Loop. An organization will Observe and Orient through these facets to develop situational awareness, and then Decide and Act through these same facets in order to achieve the desired effect. Action through these facets with respect to the OODA Loop must be consistent with defined cybersecurity governance goals. Through successive OODA Loops, action must be consistent with the mission of the organization and direction provided through sources such as statutes, orders, regulations and other guidance.

³ Although the maximum technology-related decision is limited to approximately three years due to rate of technological change, government agencies must program their expected budget needs five years in advance. In addition, the DoD is legislatively mandated to formulate strategy and priorities through the Quadrennial Defense Review process.

⁴ These nine facets are derived, in part, from the Joint Capabilities Integration and Development System (JCIDS) Manual, 19 January 2012. They reflect Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy (DOTMLPF-P) Analysis, which is part of Capability-Based Assessments. They are augmented by several topics described in the CERT Resilience Management Model.

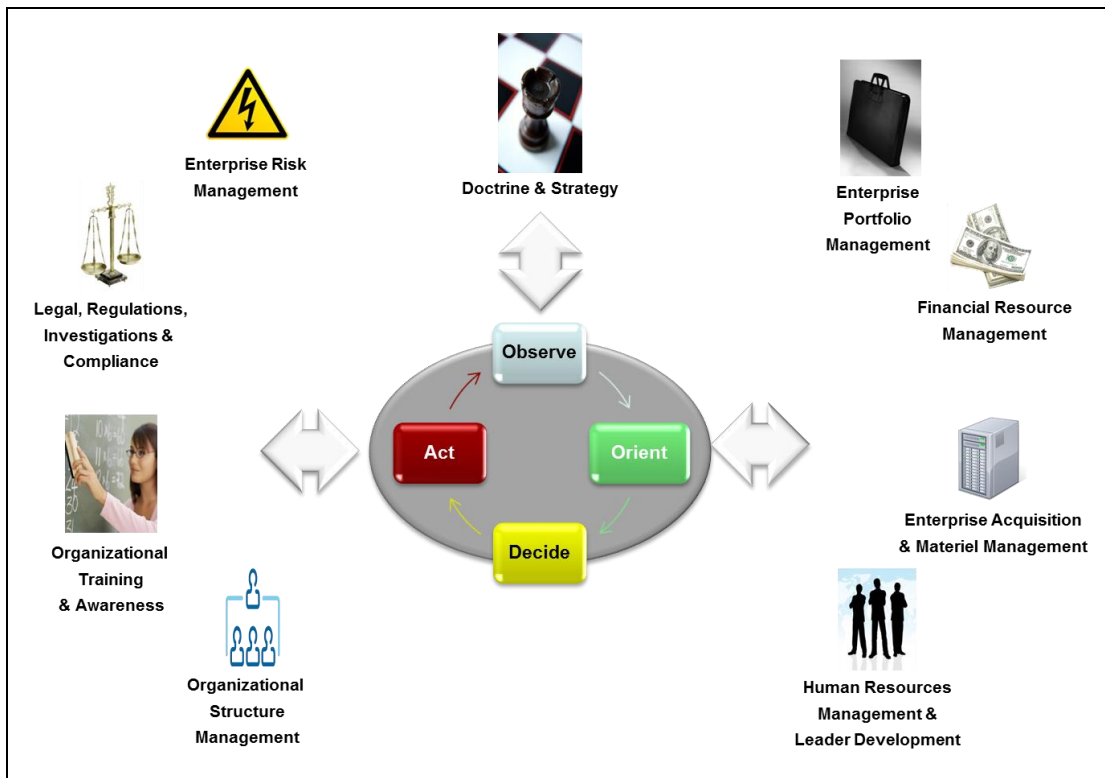


Figure 6: Facets of Cybersecurity Governance

For each of these facets, we provide a brief description and then one key goal and one or more key questions as an illustration of essential steps that can aid in defining meaningful metrics at the governance level.

3.2.1 Doctrine and Strategy

Merriam-Webster defines *doctrine* both as “something that is taught” and “a statement of fundamental government policy especially in international relations.” For the purposes of this report, a cybersecurity governance doctrine is a body of documentation outlining a common philosophy for government agencies to follow with respect to cybersecurity governance. A common doctrine allows cybersecurity decision makers across the federal space a platform from which decisions can be derived.

A strategy takes that doctrine and applies it to goals to be achieved during certain time frames. This strategy would assign responsibility and would articulate the benefits of achieving those goals. One example goal for doctrine and strategy might be

There is sufficient strategic guidance from federal government senior stakeholders to inform agency cybersecurity strategies, objectives, and key performance indicators.

Supporting questions that further elaborate this goal might be

What key guidance is missing, the absence of which is creating significant risk at the agency level? What insights at the agency governance level can be used to fill gaps at the federal senior stakeholder level?

3.2.2 Enterprise Portfolio Management

Oversight of the enterprise portfolio, particularly as it relates to IT and cybersecurity, is the application of systematic management techniques, to determine valuation of (for example, cost/benefit analysis), and performance measurement of, IT and cybersecurity planned initiatives, projects, services, infrastructures, and applications. This oversight includes determining what to continue to invest in versus what assets to retire. One example goal for this facet might be

Investments in portfolio assets are optimally balanced to achieve the mission of the federal government/agency and minimize exposure to high-impact risks. (Optimally balanced connotes that there is a set of defined criteria.)

Supporting questions that further elaborate this goal might be

Are we satisfied that we are investing in the right mix of portfolio elements (currently, one year, three year, and five year)? If not, what corrections are required, by whom, and by when?

3.2.3 Financial Resource Management

The purpose of financial resource management is to request, receive, manage, and apply financial resources (funding) to support cybersecurity objectives and requirements. One example goal for this facet might be

The allocation of financial resources reflects the considerations and priorities for all nine facets of cybersecurity (or other articulation of the scope of cybersecurity governance).

Supporting questions that further elaborate this goal might be

Are we satisfied that we have properly allocated resources for cybersecurity governance? If not, what corrections are required, by whom, why, and when?

3.2.4 Enterprise Acquisition and Materiel Management

This facet includes actions necessary to acquire, equip, operate, maintain, and support all cybersecurity activities, including all facility, information, and technology assets. One example goal for enterprise acquisition and materiel management might be

The acquisition and materiel management life cycles adequately reflect cybersecurity requirements and reviews at key decision points.

A supporting question that further elaborates this goal might be

Are the operational risks that emerge when cybersecurity is not considered during these life cycles sufficient to warrant changes in the processes used to execute these life cycles?

3.2.5 Human Resources Management and Leader Development

The development of all human resources, including federal government leaders, encompasses the employment life cycle and the measurement of staff performance in a manner that contributes to the government's cybersecurity posture. This facet also ensures that qualified personnel exist to

support all cybersecurity operations. One example goal for human resources management and leader development might be

Agency leaders and cybersecurity leaders possess the demonstrable skills to make cybersecurity governance decisions, ensure that the necessary actions are taken, monitor performance, and course-correct when necessary.

A supporting question that further elaborates this goal might be

How is performance measured for each leadership role? How are the risks arising from performance gaps identified and mitigated?

3.2.6 Organizational Structure Management

Organizational structure provides the means by which individuals cooperate systematically to accomplish a common mission. It also describes roles and responsibilities at each level of the structure. One example goal for organizational structure management might be

There is a clear, documented assignment and definition of decision-making roles, responsibilities, and authorities.

A supporting question that further elaborates this goal might be

Where is the absence of such clarity prohibiting our ability to make effective, timely decisions or creating barriers to the same?

3.2.7 Organizational Training and Awareness

Organizational training and awareness ensures that all government and service provider personnel are sufficiently aware of and trained in the skills, knowledge, and abilities required to perform their cybersecurity roles and responsibilities. Training includes the use of tactics, techniques, and procedures and often includes exercises and rehearsals. One example goal for this facet might be

Staff at all levels are adequately trained and aware, commensurate with their roles and responsibilities. (Criteria are specified to define what is adequate for each role.)

Supporting questions that further elaborate this goal might be

Is the investment in staff certifications resulting in a sufficient benefit/return? If not, where should such investment be redirected?

3.2.8 Legal, Regulations, Investigations, and Compliance

This facet ensures that the federal government is aware of and complies with all laws, regulations, policies, standards, guidelines, and other obligations such as contracts and service-level agreements. Investigations are initiated as required to examine areas of non-compliance. Furthermore, this facet provides feedback into the legislative process so laws remain relevant to the cybersecurity environment. One example goal for legal, regulations, investigations, and compliance might be

The compliance of federal agencies with current laws and regulations is resulting in a measurable improvement in their security posture (based on at least two compliance reviews).

Supporting questions that further elaborates this goal might be

Is the investment in compliance reviews resulting in sufficient benefit/return compared to the cost of conducting such reviews? If not, what improvements are needed in current laws and regulations (such as the Federal Information Security Modernization Act (FISMA), for example) and/or in the review process?

3.2.9 Enterprise Risk Management

Enterprise risk management includes the identification, analysis, and mitigation of operational cybersecurity risks to government assets that could adversely affect the ability to achieve specified missions and the ability to operate and deliver government services. One example goal for enterprise risk management might be

All risks (or selected risks or categories of risk) above established thresholds have been identified, prioritized, and dispositioned (actions taken to mitigate, control, accept, or manage as residual risk, etc.).

Are we satisfied that we are managing the right risks with the right priorities? If not, what corrections are required and by whom?

4 Enabling Data-Driven Decision Making

If we were to place this framework in the context of the human body, the OODA Loop would be the skeleton providing structure to the body, while the Facets of Cybersecurity are the muscles and organs enabling the body to operate. Data then is the life blood. How the federal government collects, correlates, measures and presents this data enables the larger body to operate and thrive. For automated systems, each change of data or system state (or the lack thereof) is an opportunity to create data that can be collected, measured, put into context, and acted upon. The data resident in government systems placed in the context of the OODA Loop can form a symbiotic process through which federal decision makers can achieve awareness and attain and act upon information faster and more effectively.

Data must be collected for a reason and in a disciplined fashion. Collecting for collection's sake can create a vast trove of important data that, due to its size and inconsistency, not only fails to increase situational awareness, but can actually impair it. Finding the nuggets of gold among the sand and pyrite becomes more and more difficult. When found, the data can actually inhibit comparisons because two collections targeting the same type of data might not itemize data in comparable ways. Although there are mathematical constructs to compensate for these differences, there is no substitute for collecting data with an eye toward apples-to-apples comparison. In the following sections, we discuss an overview of data engineering and the types of cybersecurity governance-related data that should be targeted. We also discuss ways to process and present the data to derive useful meaning for cybersecurity governance decision makers. Finally, we discuss how to leverage this information for optimal governance decision making and execution.

4.1 Collecting Situational Awareness Data and Information (Observe)

In order to improve cybersecurity governance across the federal space, we must first identify and collect the data that will be used over the rest of the OODA Loop. At the core of cybersecurity governance is the Enterprise Risk Management cybersecurity governance facet. By targeting data functionally, we can ensure a greater consistency of meaning and a greater utility for cybersecurity governance purposes. For the purposes of clarity, the following conventions are used with respect to risk management. (This list does not seek to dispute other uses of these terms used in other reports and by other organizations.)

- A **Threat** is used to denote an active entity who purposely or incidentally exploits a cybersecurity vulnerability. The threat can be a person, software, or an environmental occurrence. A threat is evaluated by its willingness and ability to exploit a vulnerability resident in defended assets. Measurement of threat-related factors is discussed in Section 4.1.4.
- A **Vulnerability** is used to denote a condition that enables a threat to degrade the confidentiality, integrity, or availability of an organization's information assets in order to create an impact. Factors such as a vulnerability's prevalence on the network (both in quantity and age) and the exploitability of the vulnerability support measurement of vulnerability-related factors. Measurement of vulnerability-related factors is discussed in Section 4.1.5.
- An **Impact** is used to denote a degradation in an organization's ability to execute its mission-essential functions. Measuring impact typically involves subjective, qualitative analysis

based on perceived mission impacts of the risk. Measurement of risk impact takes place when one is considering threats and vulnerabilities that can intersect at a certain asset, group of assets, or organizational mission.

Data-driven cybersecurity governance must be accomplished as a data engineering process that is integrated with data engineering taking place to support operational security. As discussed in Section 3, with governance we are more concerned with aggregations and trends. However, a symbiotic relationship must exist between data that is collected for operational cybersecurity purposes and that is collected, processed, presented, and acted upon for cybersecurity governance purposes. Decoupling data management at the governance and operational levels can lead to decision makers at the two levels becoming unsynchronized. This report therefore provides a discussion of operational data collection. In Section 3, we will discuss how this data can be put through various quantitative techniques to provide governance meaning.

In order to set risk-related data in the governance context, management-related information must also be collected. This data is discussed further in above and enables the cybersecurity governance decision maker to approach the data from the perspective of the Facets of Cybersecurity governance discussed in Section 3.

4.1.1 Authoritative vs. Non-Authoritative Data Sources

Authoritative data sources have a major input in implementing cybersecurity governance. Data sources are identified as authoritative based on their ability to stand alone as a source for one or more facets of cybersecurity governance. Data is being collected in central repositories from federal D/As in order to provide visibility into areas such as federal workforce, performance, spending, and cybersecurity risks.

Authoritative data sources are those whose population and breadth of criteria are such that they can be considered acceptably descriptive of outcomes within a certain measurement category. For instance, data collected from the Vulnerability Management capability of the Continuous Diagnostics and Mitigation (CDM) program or the Policy Auditor component of the DoD Host-Based Security System's (HBSS) can be assumed to be descriptive of an organization's vulnerability management outcomes. However, non-authoritative data sources, such as DHS's Risk and Vulnerability Assessments (RVA) or the DoD's Command Cyber Readiness Inspection (CCRI), may be such that they do not cover enough of the population or enough criteria of the category being evaluated to be authoritative. Their value, however, can reinforce or reduce the level of confidence in the authoritative data source, depending on whether they tend to agree or disagree with the authoritative data source.

It is important to understand that a data source is authoritative only if it is intended to be authoritative. If the quality of the data is insufficient, then including it in calculations should be a forcing agent to improving data quality. Governance practitioners should seek to improve data collection prior to abandoning an authoritative data source, since the same conditions that led the authoritative data source to exhibit poor data quality can very likely be reproduced were the source to be replaced.

Since this framework focuses on strategic-level governance, sources such as the Office of Personnel Management's (OPM) FedScope and the Office of Management and Budget's (OMB) IT

Dashboards act as additional authoritative data sources to help measure whether an outcome is efficient given the resources applied.

4.1.2 Data Engineering

Before delving into the major categories of targeted data, a brief discussion of data engineering is required. Data engineering represents the body of effort surrounding systems and infrastructure to perform data collection, organization, and management. In the modern world of Big Data, data engineering involves designing, implementing, synthesizing, optimizing, and maintaining systems that enable the storage, access, and manipulation of data. The four Vs of Big Data (volume, velocity, variety, and veracity) represent high-level target metrics that data systems must be engineered to meet. Data engineers must design systems to store and provide access to massive volumes of data, often in a large variety of formats and structures. These systems often must be capable of ingesting data and performing analyses in real time. Finally, systems must be able to maintain the veracity of data stored, ensuring data integrity and availability throughout any system of manipulations and transformations.

While the finer points of data engineering are beyond the scope of this report, due to the ever-evolving architectural nature of the subject, the first step in developing a data-driven cybersecurity solution is developing a methodical data-engineering approach. A disciplined data-engineering approach not only focuses on the people, processes, and technologies within the control of the measuring organization itself but also the interdependencies inherent in integrating data from various heterogeneous end points. Collection of target data sets should be streamlined so that collection is performed at the minimum number of locations. Duplicative collection not only draws on finite organizational resources but also decreases the validity of the data as the meaning of the redundant data streams will likely diverge. Reporting organizations may find themselves having to choose between collection requirements to the detriment of all the collection requirements. Two approaches to consider are

- standardize collection guidance at the highest level, while leaving collection planning and execution to subordinate organizations
- leverage a centralized collection system, while logically separating the data through access controls, so that two organizations can use the same resources while retaining control of their data

4.1.3 Management Information

Cybersecurity governance is very often affected by variables outside of cybersecurity itself. Therefore it is important that the cybersecurity governance decision maker have a situational awareness of those facets that affect cybersecurity. These are represented by the Facets of Cybersecurity Governance described in Section 3.2 of this report. Maintaining a good working relationship with the custodians of data in those facets can provide insight about forces that affect cybersecurity governance. By developing goal-based metrics, the cybersecurity decision maker can identify these trends and take well-researched, disciplined action in those areas. Sources of this data include

- human resource databases
- funding and spending databases

- training databases
- organizational structure databases
- databases of existing, prior, and prospective contracts
- property databases
- knowledge bases of related statutes, executive orders, and other mandates

Collection and usage of this data is an important differentiator between support of governance and operational OODA Loops. Much of what has been described in the preceding portions of this section has been operational in nature. By correlating operational cybersecurity data with data pertaining to other Facets of Cybersecurity Governance, enterprise-level decision makers can hone in on key areas for improvement that can have the greatest effect across the government. For instance, a metric describing the number of incidents reported by an organization, by itself, provides little actionable value. However, if analyzed per capita, it might reveal that the organization is reporting at a rate far below what could be expected from an organization of its size. This might indicate a lack of effective incident reporting procedures. If the data can be further correlated to the percentage of series-2210 vacancies, it might indicate that high rates of vacancies or personnel turnover is inhibiting incident reporting.

Often the organizations that manage this data publish it, such as in the case of OPM's FedScope employment cubes and OMB's IT Dashboard. By defining use-case-based metrics, a cybersecurity governance organization can coordinate system interface agreements where routinely queried data can be provided on a recurring basis.

4.1.4 Threat Information

Today, much of unclassified information regarding threat actors, their prevailing attack patterns, and their preferred targets are locked up in disparate incident reports, court records [Glenny 2011 Kindle location 4230] books, news stories, and database entries. Freelance threat actors themselves can be an excellent source of information either through anonymous chat boards or through direct interaction [Glenny 2011 Kindle location 4244]. Through techniques such as text analysis and machine learning, patterns can be derived that can not only reveal the prevailing attack patterns of certain threat groups, but also their preferred target types.

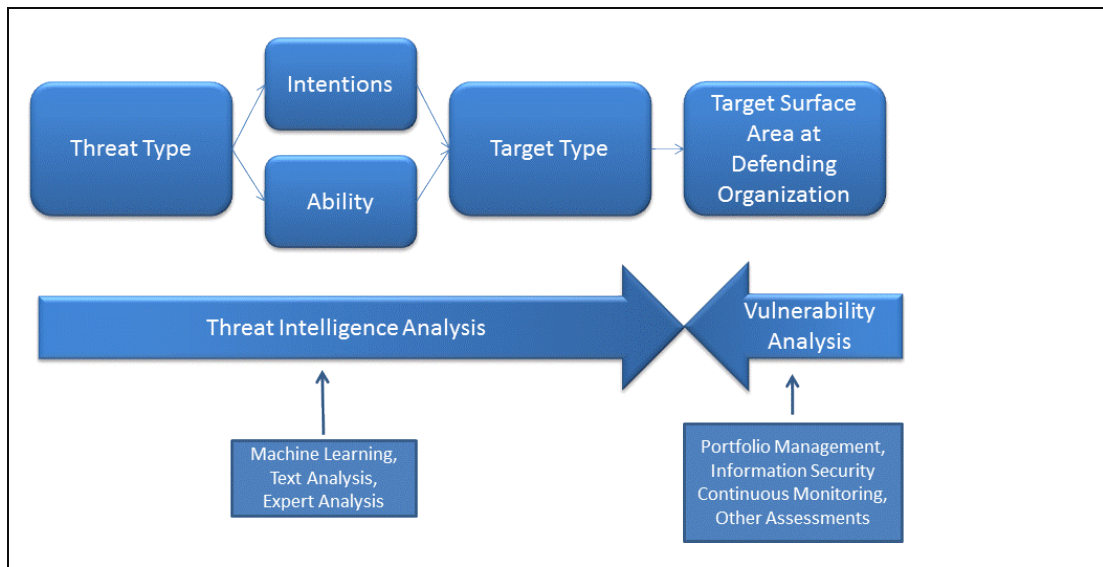


Figure 7: Matching Prevailing Attack Patterns to an Organization's Target Surface Area

By comparing this pattern information with the automation assets being defended, cybersecurity decision makers can fine-tune their use of cybersecurity governance facets and enablers to prioritize their strategy. Policies, training, and awareness practices can be made more agile to account for these trends. Acquisition strategies can be developed to better target finite resources to the attack patterns of greatest probability and potential impact to that organization. By communicating these relationships in language accessible to non-technical leaders, cybersecurity governance decision makers increase the chance of garnering the leadership support necessary to protect the assets key to the organization's mission success.

This is not static. This relationship must be viewed as dynamic, with cybersecurity governance decision makers continually reassessing their priorities and strategies. The faster the governance decision makers can assess changes in the threat landscape, the sooner they can initiate changes that will often have long lead times before operational change is achieved.

Internal threats can have just as much or more impact as external ones. Similarly, physical threats must also be incorporated in any listing of potential threats to an organization. In many cases, threats can fall into more than one category. Blended threats are the intersection of adjacent threat types. For example, an internal employee may, for financial gain, work together with an external threat to exploit security.

More information on collection and analysis of threat information can be found in Appendix A.

4.1.5 Automated Vulnerability Information

4.1.5.1 Challenges in vulnerability information collection

Due to the complexity of government networks and the rapid discovery of vulnerabilities, our traditional, legacy, manual methods to reduce vulnerabilities are no longer sufficient. The one central, consistent theme that has emerged throughout modern risk management methodologies and programs such as Information Security Continuous Monitoring (ISCM) has been that automated

vulnerability management is paramount to keeping pace with the dynamic nature of today's threats.

Traditionally, vulnerabilities have been synonymous with software weaknesses (such as those leading to memory buffer overflows or denial of service attacks). However, other vulnerabilities and weaknesses exist in an organization that act as platforms to obtain, update, or delete potentially sensitive information. These include weaknesses or vulnerabilities in how an organization has deployed and administered

- hardware asset management (including removable media)
- logical and physical privileges and authorizations
- credentials and authentication
- security clearances and suitability
- network architectures and topologies

Organizations are starting to explore the opportunities to automate vulnerability identification. The ability to migrate (as much as possible) to an automated state of vulnerability management is critical to an organization's ability to maintain an acceptable security posture.

More information on collection and analysis of vulnerability information can be found in Appendix B.

4.2 Position the Enterprise for Action (Orient)

Next in the OODA Loop comes the point at which cybersecurity decision makers make sense of the data they've collected in relationship to other data sources and the larger environment. In order for the federal enterprise to gain and make sense of situational awareness, it is important for the OODA Loop to be responsive to how large governmental organizations function. In order to better synchronize the efforts of such a large, heterogeneous organization, a common understanding must be developed of cybersecurity governance expectations. This enables governance decision makers to not only see data in the same way but to have a more unified understanding of what the desired outcomes should be.

Many observe government through the fallacy that large organizations act as one large entity with a unified set of intents and desired outcomes. Graham Allison and Philip Zelikow termed this the Rational Actor model of government decision making in their book, *Essence of Decision: Explaining the Cuban Missile Crisis* [Allison 1999, Kindle location 508]. However, in order to develop a common situational awareness and a coordinated approach to acting upon that awareness, more nuanced models of thought are required. These more nuanced models enable those who collect, analyze and make decisions upon data to do so with a more effective understanding of those affected and who affect the process.

One example is the Organizational Behavior model. Through this model, government actions are viewed as outputs developed as a function of strategies, processes and procedures developed in advance of the output [Allison 1999, Kindle location 3235].

To perform complex tasks, the behavior of large numbers of individuals must be coordinated. Coordination requires standard operating procedures: rules according to which things are

done. Reliable performance of action that depends upon the behavior of hundreds of persons requires established “programs” ... [Allison 1999, Kindle location 3240]

The behavior of these organizations—and consequently of the government—relevant to an issue in any particular instance is, therefore, determined primarily by routines established prior to that instance. Explanation of a government action starts from this baseline, noting incremental deviations. But organizations do change. Learning occurs gradually, over time. [Allison 1999, Kindle location 3246]

By developing flexible OODA-based processes that are widely understood and can adjust to changes in the environment, the federal government can make faster and more effective cybersecurity governance decisions. These decisions can be, in turn, executed more effectively, and can lead to the perpetuation of faster, more effective follow-on OODA Loops.

To be effective, metrics and metric-driven products can derive greater impact through tailoring to affected stakeholders. This is done through another model, which Graham and Zelikow called the Governmental Politics model.

Both by charter and in practice, most players “represent” a department or agency along with the interests and constituencies their organization services. Because their preferences and beliefs are related to the different organizations they represent, their analyses yield conflicting recommendations. Separate responsibilities laid on the shoulders of distinct individuals encourage differences in what each sees and judges to be important. [Allison 1999, kindle location 5603]

“Politics” in this regard does not indicate a self-serving nature, but rather the differences in viewpoint that each decision maker’s mission engenders. These differences mean that those who develop cybersecurity-governance-related products must understand these unique mission-related viewpoints and tailor outputs accordingly. Through data science and visualization, the cybersecurity governance decision maker can deliver the products necessary for each decision maker to perform according to the facets of cybersecurity more relevant to him or her.

4.2.1 Defining Environmental Constraints for Implementation of the OODA Loop

4.2.1.1 Doctrine and Strategy

Doctrine and strategy, as described in Section 3.2.1, provide an important frame of reference for the Orient phase of the OODA Loop. Depending on how it is used by the issuing organization, some or all of it may not be mandatory. However, it provides a common understanding of what the desired methods and outcomes are across disparate decision makers. Given the relative newness of the cybersecurity governance field, doctrine and strategy therein has been in a state of flux and development. When assessing doctrine and strategy during the Orient phase, even shortcomings in doctrine and strategy can be important information to consider. Greg Wilshusen, a director in GAO’s Information Technology team, stated the following in a GAO Podcast dated February 14, 2013:

While the government has developed various strategy-related documents over the years that address aspects of these challenges, it has not yet developed an overarching cybersecurity

strategy that articulates policy actions, assigns responsibilities for performing them, and establishes time frames for their implementation. We are recommending that the White House Cyber Security Coordinator develop a federal cybersecurity strategy that includes all the key elements of...desirable characteristics of a national strategy. This strategy should also be used to better [ensure] that federal government departments and agencies are held accountable for making significant improvements in cybersecurity challenge areas by, among other things, clarifying how oversight will be affected. [Wilshusen 2013].

4.2.1.2 Constraints and Mandates

Cybersecurity governance decision makers operate within a hierarchically constrained environment. Chief among these constraints are statutes laid down by Congress. These statutes are often broad in nature, allowing for the details of execution to be outlined by elements of the executive branch. An example of this is FISMA. FISMA directs the director of the Office of Management and Budget to oversee agency information security policies and practices [U.S. Congress 2014]. Figure 8 shows the hierarchy of mandates and constraints that must be considered in order orient analysis and decision making by cybersecurity governance decision makers.

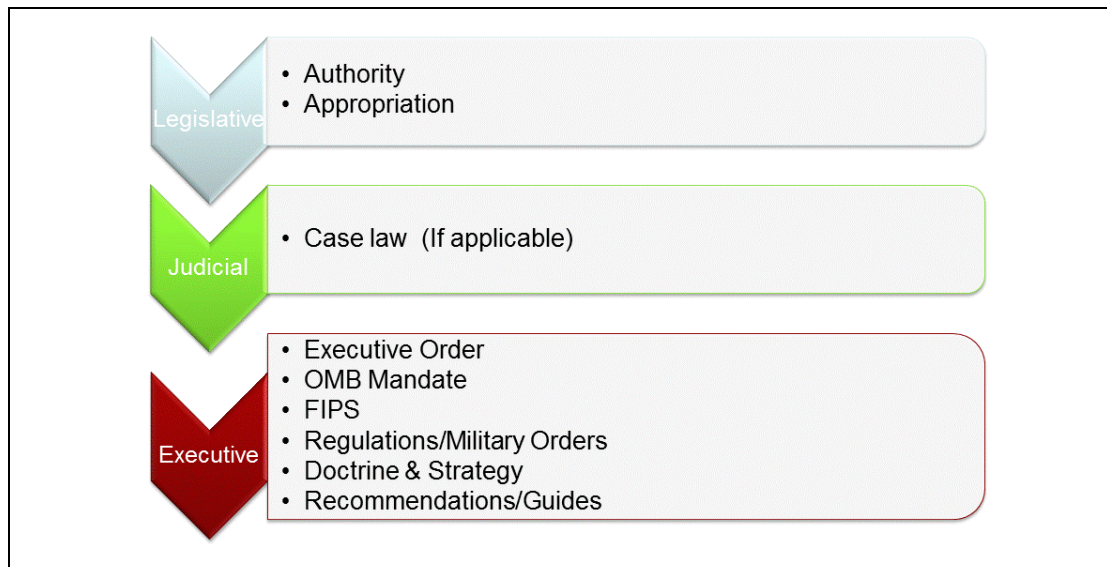


Figure 8: Sources of Mandates and Constraints

Once enacted, some statutes may be challenged or referenced in court cases. In such cases, the findings of the courts may change or clarify how those statutes are carried out. (Different courts can even have differing, even contradictory, interpretations of the law.) When this happens, the applicable case law can become a consideration in how the law is carried out. Cybersecurity governance decision makers should therefore maintain regular contact with legal advisors when determining how a statute affects a particular OODA Loop.

Laws in turn are most often amplified by executive order, OMB guidance, and regulations. These documents provide mandatory executive guidance on how the authorities outlined in the statute are to be executed. They are key to all facets of the OODA Loop. First, they provide a basis for determining what data must be collected in order to satisfy mandates that they contain. Second, they enable organizations to develop benchmarks against which the data can be compared and an

indication as to the stakeholders most likely to require that information. By developing benchmarks, the cybersecurity governance decision maker can ensure that situational awareness supports the development of the right decision at the right time. Finally, it enables execution to be measured against the intended outcomes delineated in the mandates.

Although generally not mandatory, guides and recommendations, such as NIST SPs and military field manuals, can provide context and expert insight that can make metric-based products more effective.

Constraints and mandates provide an important input in putting data collected in the Observe phase into context. They often drive the reasoning behind developing expected states of the data, such as benchmarks and targets. When the data deviates from these benchmarks and targets, they can be a starting point for root cause analysis.

4.2.2 Data Science

After setting the context for data collected in the Observe phase, the development of a data science approach enables the cybersecurity decision maker to make sense of the data and turn it into information. Data scientists are often adept statisticians and machine-learning experts who focus on designing algorithms and analytical techniques to answer challenging questions from data sets. The primary goal of data science is to turn data gathered into actionable knowledge or insights for an organization. Data scientists must be capable of employing cutting edge technologies to analyze data in a variety of scenarios. For example, some problems may call for a high-powered batch processing analytics solution, while others require the speed of real-time analytics. Data scientists not only design algorithms, but often also possess expertise in data modeling, helping to design the data itself to enable the necessary analysis techniques.

Assimilating the results of data-driven analysis into indices provides metrics that are accessible and actionable for decision makers at the governance level. When properly constructed and explained, indices provide a yard stick for determining thresholds of action, as well as a starting point for root-cause analysis. Common examples of indicators include the Bureau of Labor Statistics' Unemployment Rate and Consumer Price Index. Indices can provide a one-stop indication of cybersecurity governance across government and private-sector organizations, but can also provide entrée to underlying metrics and indices. These underlying metrics and indices can bring into sharp relief causal factors that can underpin the development of hypotheses and theories. These hypotheses and theories enable cybersecurity governance decision makers to develop well-founded decisions and actions further on in the OODA Loop.

Additional information regarding how indices and other quantitative methods can be used to enhance cybersecurity governance decision making can be found in Appendix C.

4.2.3 Visualization

Visualization has proven effective to aid in understanding large, intrinsic data commonly found in large-scale scientific simulations and biomedicine [Fan 2013]. Without effective data visualization, even the best data engineering may not suffice. The challenge is to develop visual representations, layout methods, user interfaces, and interactive techniques that can effectively facilitate visual data mining, interrogation, and communication of the vast amounts of cybersecurity information to the appropriate audience. An effective visualization approach enables the following:

- faster time to problem resolution with cognitive-research driven visualization of context to discover cybersecurity governance challenges
- contextualization for maximum visibility with extensive data correlations, using batch analytics and machine learning algorithms for insightful intelligence
- increased ROI of existing security infrastructure, by bolstering the intelligence of legacy security solutions and integrating with existing security appliances and systems

The challenge is to obtain the best visual representation for each type of analysis task and the individual who will be accessing the data. Understanding the model in which the user operates is imperative (discussed in Section 4.2). Whether the user intends to use the data to fulfill a predetermined process need or to fulfill a stakeholder's information need will drive how the data should be visualized. Using summary interfaces that enable drilldown can facilitate both quick access and root-cause analysis. It is critical to maintain unbroken analysis context while drilling down into the details, so that the analyst can classify the pattern he or she sees in the data.

Many visualization systems do not get widespread adoption because they confront the user with sophisticated operations and interfaces. We suggest extending the visualization systems with a learning capability to improve both their performance and usability. This can be done by including volume segmentation, flow feature extraction, and clustering, to illustrate how machine learning can help streamline the process of visualization, simplify the user interface and interaction, and support collaborative work. This can be captured in a scorecard visualization that provides simplicity of issue and remediation description, as well as layered optional drilldown.

4.2.3.1 The Governance Scorecard

The improved visualization techniques used for the scorecard will enable decision makers to more readily acquire and employ the information they need to make decisions. This emphasis on visualization will focus not just on measurements themselves, but the resources a leader has to mature his or her processes and practices.

The benefits of a scorecard from other viewers and dashboards is the ability to incorporate information, derived from data and scoring algorithms, around how to enhance an individual organization's grading for cybersecurity governance improvement. The transformative roadmap that is provided internally to an organization can be achieved for the measured organization through an improvement scorecard.

The scorecard establishes a measurement approach for cybersecurity governance that can be adjusted for new technologies and can be recalibrated for changing practices and processes surrounding the technology. This fluidity between authoritative data sources feeding a scorecard measures both strategic cybersecurity governance and provides situational awareness for tactical and necessary strategic change needed.

4.2.3.2 Usability Requirements

An interface should be easy to learn how to use and easy to remember how to use. The latter pertains especially to devices that require infrequent use.

Table 3 shows the 10 usability heuristics for user interface design based on Jakob Nielsen's definition [Nielsen 1990, pp. 249-256]. We recommend using these as the basis for the system's usability and visualization requirements.

Table 3: Usability Heuristics

Show System Status	The system should always keep users informed about what is going on, through appropriate feedback within reasonable time.
Match Mental Models	Follow real-world conventions, making information appear in a natural and logical order.
Use Plain Language	The system should speak the users' language, with words, phrases and concepts familiar to the user, rather than system-oriented terms.
Prevent Errors	Careful design prevents a problem from occurring in the first place. Either eliminate error-prone conditions or check for them and present users with a confirmation option before they commit to the action.
Allow for Graceful Recovery	Users often choose system functions by mistake and will need a clearly marked "emergency exit" to leave the unwanted state without having to go through an extended dialogue. Support undo and redo and give the user control over what they are doing. Error messages should be expressed in plain language (no codes), precisely indicate the problem, and constructively suggest a solution.
Be Consistent and Use Standards	Users should not have to wonder whether different words, situations, or actions mean the same thing. Follow platform conventions.
Foster Recognition Rather than Recall	Minimize the user's memory load by making objects, actions, and options visible. The user should not have to remember information from one part of the dialogue to another. Instructions for use of the system should be visible or easily retrievable whenever appropriate.
Anticipate Needs	Accelerators—unseen by the novice user—may often speed up the interaction for the expert user such that the system can cater to both inexperienced and experienced users. Allow users to tailor frequent actions.
Design with Minimalist Aesthetic	Dialogues should not contain information that is irrelevant or rarely needed. Every extra unit of information in a dialogue competes with the relevant units of information and diminishes their relative visibility.

4.2.4 Prioritization of Problems to be Addressed

Once problems to be addressed and their root causes are identified, it is incumbent on the cybersecurity governance practitioner to prioritize the problem in relationship to other problems being considered. There are three high-level inputs that should be considered when determining how a problem should be ranked. These aggregate metrics are, in turn, fed by a hierarchical set of supporting numerical inputs. Several mathematical models can be used to organize and express these relationships into a single prioritization metric. The purpose of this model is not to deliver an unquestionable priority level for the problem, but rather a starting point for prioritization analysis. "What-If" analysis should be performed by changing the weights of inputs to this model, and the final result should be assessed qualitatively by decision makers.

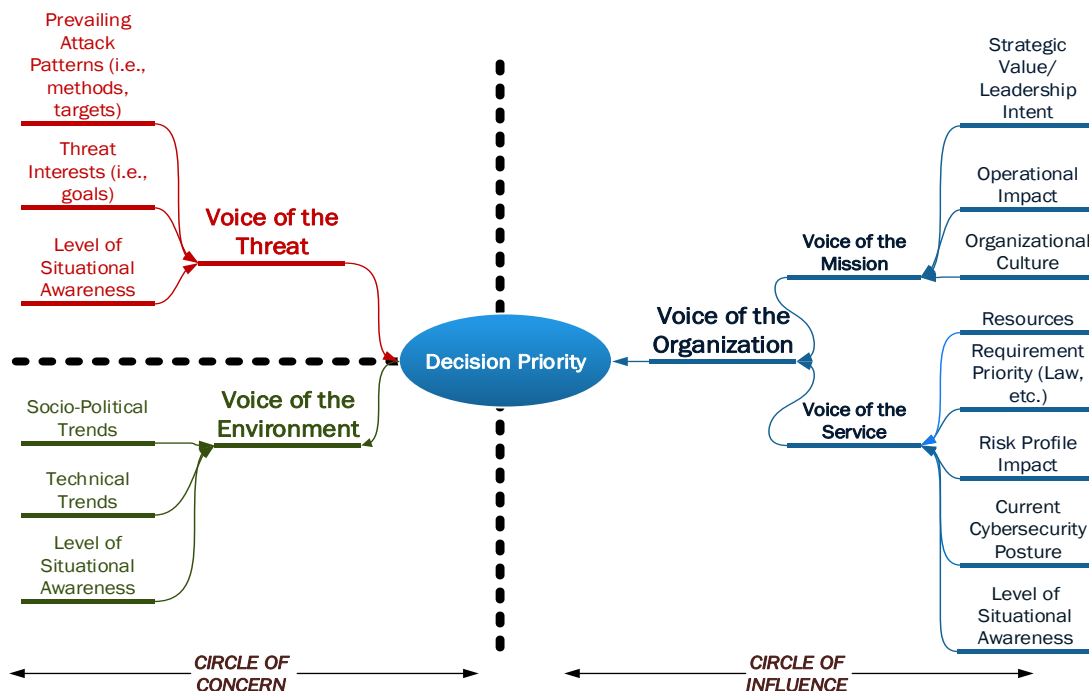


Figure 9: Problem Prioritization Decision Tree

The Voice of the Organization numerically expresses considerations from the vantage point of the defending organization. It is an outward-looking, aggregate metric that ensures that multiple internal considerations are taken into account in the prioritization. These are considerations that the organization can affect, or as author Stephen R. Covey would term, are within the organization's "Circle of Influence" [Covey 1989, p. 82].

This macro metric is in turn fed by two subordinate aggregated metrics. The Voice of the Mission considers inputs that effect the business functions and mission accomplishment of the organization. Example inputs can be a qualitative input that represents how addressing the problem supports leadership's strategic direction for the organization. This input can be generated through expert opinion and techniques such as the Delphi Method. For instance, one problem may have outsized impact on a unit or division that has a mission of key importance to leadership. Operational impact is a qualitative measure that would represent the impact (positive or negative) solving the problem would have to mission accomplishment. Organizational culture indicates how readily group behavior would support solution of the problem.

The Voice of the Service considers inputs that represent how the cybersecurity service itself is delivered. Resources indicates the degree to which current resources (i.e., people, information, technology and facilities) support or do not support problem solution. It is important to note that in cases where the root cause of a problem is a resource problem, this metric could have an inverse relationship to the priority. That is, the lower the on-hand resources, the greater the priority to fix the problem.

Along with the Voice of the Organization, the problem prioritization should also consider two other, externally focused aggregate metrics: the Voice of the Threat and the Voice of the Environment. These provide numerical input regarding matters which are outside of the organization's

control, or as Covey would term, those within the organization's "Circle of Concern" [Covey, 1989, p. 81].

The Voice of the Threat indicates how ongoing threat actor goals, target types, and methods indicate the level of risk inherent in the problem. For instance, if a certain category of threat actors have been making good use of a certain attack pattern that is relevant to the problem at hand, then the Voice of the Threat metric could cause the prioritization of the problem to increase. This data can be gleaned by techniques such as forensic trends, sentiment trends in key hacker forums, and other intelligence. By developing indices around common domain values, such as those categorizing target types, attack techniques, vulnerabilities used, and heuristic network patterns, these seemingly unmeasurable patterns can be mathematically gauged and provide more real time input on threat activities.

The Voice of the Environment measures and provides input on factors such as economic, socio-political, political-military, and cultural trends. In addition, by capturing significant types of events (i.e., initiation of military hostilities between two other nation states), we can identify whether such events had a consistent or recurring effect on problems such as the one being considered. Examples include political-military conflict between Russia and Estonia, Russia and Georgia, and Russia and Ukraine, which each led to temporary changes in the cybersecurity landscape.

Effort is required to properly tune this decision tree to most consistently deliver prioritization information desired by the organization. It is important to note that such automated decision support is intended to make for faster, more informed decisions. It does not replace the informed and educated analysis of leaders and their staffs.

4.3 Key Planning and Decision-Making Factors (Decide)

Each metric should be tied to an organizational goal. Therefore, the courses and plans of action developed in the Decide phase of the OODA should also seek to improve attainment of those goals. Focusing on simply improving the end result of a metric is to miss a key opportunity for improvement of the cybersecurity governance process.

Investment in a disciplined process in the Decide phase of the OODA Loop substantially increases the chance of success in the Act phase. It ensures that the maximum practicable set of considerations and courses of action are considered prior to committing precious organizational resources. According to Stephen Covey in his book, *The 7 Habits of Highly Effective People*

"Begin with the end in mind" is based on the principle that all things are created twice. There's a mental or first creation, and a physical or second creation to all things.

[Covey, S.R. (2013), pp. 98-99]

As much care as possible should be taken in the development of the decision to ensure it is valid and executable. It is similar to the old carpentry adage that one should "measure twice and cut once". The decision-making process should be as inclusive as possible of those who will execute in the Act phase [Deming 1986, pp. 107-108]. This idea extends not only to the organization's IT community, but also to those who will support through one of the Facets of Cybersecurity Governance, such as acquisition, legal, Human Resources, and training personnel. This inclusion will eliminate a bias toward the decision maker's own comfort zone. Leveraging trained facilitators to

elicit and organize considerations and courses of action can lead to a well-rounded, practical decision that enables execution and improves cybersecurity governance.

4.3.1 Determine Time Available for Decision Analysis

Cybersecurity governance decision makers should not take the lion's share of time available to make a decision. In a multi-echelon organization like the U.S. government, nested decision making will take place at several layers of the organization. A rule of thumb is for each layer of decision making to take only one-third of the available time to develop a decision and develop a plan of action, leaving the remaining two-thirds to subordinate organizations. The fact that governance decisions have longer time horizons will often mean that there should be a greater sense of urgency to make decisions and issue guidance. Accounting for multiple feedback loops from organizational layer to organizational layer will help to fine tune the decision and avoid risks in the schedule.

Determining the time available for the decision-making process supports determining how formal and deliberate the process should be. Although all of the steps in this section should be completed, whether they should be done in formal, discrete steps depends on the time available. Shorter lead times will cause the decision cycle to be less formal and more compressed. Longer lead times, however, should allow for a greater (but finite) level of rigor. As much as unnecessarily shortening the decision-making cycle can lead to malformed decisions, excessive contemplation can lead to "analysis paralysis."

4.3.2 Determine Hypothesis or Theory

Cybersecurity governance decision makers should understand that the results of the Observe and Orient phases of the OODA Loop may not yield information that is ready to be turned into execution through the Facets of Cybersecurity Governance. Rather these results may require additional research and analysis to develop greater accuracy and assurance. Thus, one OODA Loop may not lead to change but rather to another OODA Loop to further refine the problem set. One way to recognize the difference is to determine whether the results of the Observe and Orient phases is a hypothesis or a theory.

Merriam-Webster defines a hypothesis as "a tentative assumption made in order to draw out and test its logical or empirical consequences." As discussed in the "Data Science" section of this report, correlation does not always mean causality. Although trends in data may point to a particular outcome, an absence of analytic rigor can lead to decisions and actions that needlessly waste organizational resources. A disproven hypothesis does not mean failure. The cybersecurity decision maker should test hypotheses derived from the Observe and Orient phases by performing additional analysis through a follow-on OODA Loop until there is sufficient confidence that the problem statement is accurate. In this case, deciding how the follow-on OODA Loop should be conducted would be the focus of the Decide phase.

Once sufficient confidence has been reached, a theory has been developed. Merriam-Webster defines a theory as "a plausible or scientifically acceptable general principle or body of principles offered to explain phenomena." The theory should include the identification of a root cause of the phenomenon observed. It is these root causes that the decision-making process will attempt to affect through the Facets of Cybersecurity Governance. When assessing the theory, a desired state

must be identified so that those who will carry out the decision have a clear understanding of what success looks like.

4.3.3 Determine Enablers

An enabler is a mechanism through which the cybersecurity governance decision maker will attempt to affect the problem's root cause. An enabler can be a yet-to-be developed training program in conjunction with a vehicle for delivering that training, such as an organization or acquisition vehicle. Existing legislation and policies can also be enablers, as authorities and appropriations may already exist to deliver the solution.

An understanding of the enabler's capability is critical. A statute, for instance, may appear to provide the authorities necessary to underpin a solution. However, subsequent court cases may have altered how the statute can be implemented, or whether the statute can be implemented at all. (It is possible for courts to rule that a statute is unenforceable, often due to a perceived lack of clarity in the way the law was written.) Likewise, an organization may have been given the mission to deliver a service either through legislation or delegation of authority, but a lack of resources could restrict its ability to carry out those authorities. By taking these mechanisms into account when developing the decision and plan of action, compensating enablers can be identified and execution of the OODA Loop can be made faster and more effective.

4.3.4 Determine Criteria and Weighting

Criteria are determined in order to evaluate courses of action (COA) for final selection. A key influencer in developing criteria is the goal associated to the metric under review (Section 3.2 discusses goal development with respect to the Facets of Cybersecurity). Again, improvement of the metric result should be considered a by-product of improvement, not a criterion in and of itself. The two types of criteria are screening criteria and evaluation criteria.

A screening criterion is a condition that the COA must meet in order to be acceptable. For instance, a screening criterion might be that the solution must not require funding over and above current budgeted spending levels. In this case, if a COA calls for additional funding, it is discarded and no longer considered.

Evaluation criteria are ones against which each COA is assessed with respect to other surviving COAs. These evaluation criteria can be weighted so that one is considered more significant than others. While the budgetary screening criterion above may screen out unacceptable COAs, cost can still be used to assess surviving COAs. Cost can be considered the most important consideration over, say, the technical efficacy of the COA. This is the case in acquisition actions categorized as "Lowest Price, Technically Acceptable," where the cost of a solution is considered before the technical merits. If two solutions are believed to cost the same, then other evaluation criteria can break the tie. Although there are differing philosophies on developing evaluation criteria, duplication should be eliminated as much as possible and the set of criteria kept to a manageable set. Using too many evaluation criteria can cause confusion in justifying the end decision, making it more difficult to implement.

4.3.5 Determine Courses of Action

As with choosing evaluation criteria, COAs to be considered should be kept to a manageable and feasible set. A common practice is to include “throw-away” COAs that are intended to make other COAs look more attractive to decision makers. This practice should be discouraged as causing needless distractions that can consume unnecessary analysis and reduce the decision makers’ confidence in the process. Although each COA does not need to be developed into a full execution plan initially, enough implementation detail should be developed to facilitate the evaluation of surviving COAs against the evaluation criteria.

Again, the target of each COA should not be the improvement of a metric, but rather improving the conditions surrounding the root cause of the metric value [Deming 1986, pp 70-76]. If the underlying cybersecurity posture has been improved, then it can be reasonably expected that the improvement will present itself in the metric. Target values may be set as benchmarks, but only based on solid historical data and only as a starting point for additional analysis and leadership attention if the metric value fails to meet the benchmark. Each metric seeks to measure attainment of a goal. The cybersecurity governance decision maker should be able to map the COA back to the goal at hand.

4.3.6 Evaluate Courses of Action

Each surviving COA should be evaluated according to each criterion. One way to do this is through simple comparison, for example

Criteria: Expected effectiveness: $COA\ 2 > (is\ favored\ to)\ COA\ 1 > COA\ 3$

Criteria: Speed of implementation. $COA\ 2 > COA\ 1 > COA\ 3$

Criteria: Cost. $COA\ 1 > COA\ 2 > COA\ 3$

This evaluation can also be done via a decision matrix as shown in Table 4. The decision matrix can be used to compare relative values (i.e., first, second, third most favored) or can be used to compare actual data (i.e., cost or time to implement) or a combination of the two. While a full discussion of the development of a decision matrix is outside the scope of this report, it provides a useful tool to coalesce decision-making thinking around COAs and their related evaluation criteria. It should be pointed out, however, that a decision matrix should be used to augment, not replace, experienced leadership judgment in the development of a decision. That is, weights can be increased or decreased to represent the priorities of the decision maker. However, such adjustments should not be made simply to bolster a “favorite” COA.

Table 4: Decision Matrix

		Expected Effectiveness	Time to Implement	Cost	Total	
	Weight	3	2	1		
COA 1	Raw	2	3	1	13	
	Weighted	6	6	1		
COA 2	Raw	2	3	1	13	
	Weighted	6	6	1		
COA 3	Raw	3	2	1	14	<- Favored COA
	Weighted	9	4	1		

Once a COA is selected, an execution plan can be developed leveraging selected enablers to achieve the desired effect. When developing the COA into a full execution plan, cybersecurity decision makers should ensure that enablers are utilized at the proper levels of the organization to maximize their efficacy and to account for diverse conditions that may exist at the point of execution. While the various management techniques are beyond the scope of this document, cybersecurity decision makers should bear in mind that while execution guidance can be general in nature, reporting should be standardized to the maximum extent possible so that apples-to-apples comparisons can be made between present and historical measurements and measurements between such items as organizations, networks, processes, software, and hardware.

4.4 Enabling Success at the Point of Execution (Act)

The purpose of data-driven cybersecurity governance is not merely improvement of a metric's value, but an improvement in actual cybersecurity. To accomplish this, the Act phase of the OODA Loop interacts with all phases of successive OODA Loops. Measurement points along the Act phase should be identified to ensure execution is taking place in accordance with identified goals. Decision points can be set at which time alterations to the execution plan can be made in a controlled methodical manner.

If the results of the Observe and Orient phases is a hypothesis (as shown in Figure 12), the results of the Decide phase will determine how new information will be collected and through what enablers. They will also determine the parameters through which the new OODA Loop will Orient. The new OODA Loop builds upon the previous loop to test the hypothesis and refine the decision maker's understanding of the problem. In contrast to an action based on a theory, enablers are those people, processes, and technologies best able to exact the quantitative and qualitative information necessary to facilitate the new OODA Loop. For instance, if the hypothesis were that the root cause was human-resource based, then the enabler might be the custodian of human-resource data such as the vacancy rate of 2210-series federal employee positions titled Information Security.




Figure 10: Transition from Decide to Act Based on a Hypothesis

Actions based upon theories are normally those intended to create a cybersecurity governance effect in accordance with an identified goal, as shown in Figure 13. In such cases, the governance decision maker would engage the enabler most suited to achieve that effect. Effecting change across large government organizations is often difficult. Understanding the governance model (See Section 3.2) in which the enabler is acting is crucial to ensuring desired outcomes are achieved. That is, as discussed in the Orient section of this report, an enabler operating within the organizational behavior model can be expected to operate within the standing operating procedures, routines, and outputs set for them. Whenever possible, interacting with the enabler within the confines of those routines, even if with minor adjustments, can still achieve the desired effect without undue turbulence.

However, when those procedures, routines, and outputs are insufficient to achieve the desired effect, coordination with those operating within the governmental politics model may better achieve the desired outcome. Here again, the governance decision maker can expect certain behaviors according to that model. These enablers will perceive and operate according to their roles and constituencies. Engaging them in a way that shows a value added to those roles and constituencies can reap benefits and create a momentum that will pull the action to completion. When interaction between organizations is expected to be routine, then memoranda of agreement outlining the expectations, roles, and responsibilities can facilitate faster, more effective execution.




Figure 11: Transition from Decide to Act Based on a Theory

Key to the success of the Act phase is documentation. All throughout the Act phase, a historical record should be kept outlining the actions taken, and most important, the lessons learned along the way. An efficient and effective knowledge management process is key to ensuring that each OODA Loop leverages and builds on those taking place before it.

5 Implementation Using a Maturity Model

The iterative improvement envisioned for cybersecurity governance may be best achieved through use of a maturity model. A maturity model is defined as “a set of characteristics, attributes, indicators, or patterns that represent progression and achievement in a particular domain or discipline” [Caralli 2012]. The model comprises best practices and possibly “standards or other codes of practice” of the domain or discipline. Maturity models have levels with associated attributes that identify measurable transitions from one level to another. These transitions allow the organization to “define its current state, determine its future, more ‘mature’ state, and identify the attributes it must attain to reach that future state” [Caralli 2012].

The *Framework for Improving Critical Infrastructure Security* [NIST 2014], while not technically a maturity model, displays similar features. The Cybersecurity Framework Core Functions, Categories, Subcategories, and Tiers captured within a Framework Profile can demonstrate an organization’s existing or ideal state to identify cybersecurity risk management gaps. The approach described in this report enables cybersecurity governance using the same kinds of information and decision flows among the operations, business, and executive levels as described in the Cybersecurity Framework guidance [NIST 2014, p. 12].

The transformative roadmap that is provided to an organization through a maturity model could be achieved through the cybersecurity governance scorecard described in Section 4.2.3.1. The scorecard could provide incremental benchmarks at the maturity levels and serve as a measurement of progression.

More thorough guidance for using a maturity model to advance cybersecurity governance is beyond the scope of this technical report but may be developed in future work.

6 Conclusion

The federal government is at a crossroads with respect to cybersecurity. The risks facing the federal enterprise evolve at a pace faster than current, disjointed methods of enterprise-level decision making can match. However, through a deliberate OODA-Loop-based process of collecting data, making sense of that data, and making constructive, informed decisions and controlled actions, the federal government can utilize its advantages. These advantages include economies of scale, opportunities for information sharing and access to a wide portfolio of capabilities that can be leveraged if situational awareness can be improved. Effective collection and management of data and methods to analyze and process the data can enable faster and more effective convergence of situational awareness. Decision-making processes based on this situational awareness can ensure that the federal government's cybersecurity governance capabilities are aligned to deliver maximum effect. Finally, ensuring disciplined execution can both enable success at the point of execution and further situational awareness.

Appendix A: Collecting and Categorizing Threat Information

A.1 Threats and their Potential Impact on Networks

Table 5 identifies a sample of threat attack patterns associated with successful threat attack or exploitation activities. Capturing, analyzing and reporting data regarding these attack patterns can, over time, provide valuable insight as to the intentions and abilities of certain threat categories.

Table 5: Mapping of Components, Attributes, and Desired Threat Capabilities

Network Component	Attributes (Example)	Attack Patterns that Can Be Measured Over Time
Services	<ul style="list-style-type: none">• Email• Web• Authentication• Database	<ul style="list-style-type: none">• Send email with return address of the compromised user• Read emails of network users• Delete old emails• Modify website contents• Modify database contents
Users	<ul style="list-style-type: none">• Personally Identifiable Information (PII)• Protected Health Information (PHI)• Beneficiaries (family)• Education• Military service background• Ethnicity• Organizational leadership	<ul style="list-style-type: none">• Acquire the address and phone numbers of users, middle managers, and leadership• Identify users' medical conditions• Identify the family members of network users• Identify education and military service backgrounds
Hardware	<ul style="list-style-type: none">• Router password (modify block lists)• Switch components (eliminate or add components)• Modify wireless signal	<ul style="list-style-type: none">• Reduce availability or services• Add unauthorized components• Downgrade or eliminate wireless security
Unprotected Data	<ul style="list-style-type: none">• Reports or work products not circulated outside of an organization• Source of reporting components• Schedules of delivery or integrations• Suppliers and customers	<ul style="list-style-type: none">• Identify organizational core competencies• Identify delivery cycles along with customers and suppliers
Protected data	<ul style="list-style-type: none">• Complete analysis• Proposal work• Intelligence products• Large scale PHI• Personnel rotations	<ul style="list-style-type: none">• Identify most valuable elements of an organization• Create unfair bidding environments• Identify kinetic opportunities with personnel movements

A.2 External Threat Categories

One or more threat actors that do not have authorized access to an organization's information systems are referred to as external threats. These threats use a series of tactics and techniques to gain

access and exploit security vulnerabilities for unauthorized purposes. By identifying threat categories and threat groups within those categories, a taxonomy can be created through which threat intention and ability information can be analyzed.

There are many types of external threats in both capability and desired effect. The Industrial Controls Systems Cyber Emergency Response Team (ICS-CERT) website identifies five (listed from most to least capable):

1. national governments
2. terrorists
3. spies and organized crime groups
4. hacktivists
5. hackers

The most capable external threats to an organization's information systems are sponsored by national governments. The recent indictment of five Chinese nationals, along with the Advanced Persistent Threat (APT) 1 report from Mandiant, identifies China as a nation-state threat [Mandiant 2011]. Furthermore, this report identifies APT1 threat actors as being government funded and associated with a military unit.

On the opposite end of the capability list are the hackers, who are often referred to as “script kiddies.” They use script programming and common off-the-shelf tools to attempt unauthorized access. Though their capabilities are limited, if they find access, the potential damage or loss of information can be significant.

One way to delineate between script kiddies and nation-state actors involves the resource capability of the threat. Structured threats may involve planning, custom tools, and novel techniques. Unstructured threats focus on readily available tools and simple techniques. Unstructured threats look for simple vulnerabilities, while structured threats target organizations and their assets to achieve a desired effect.

Common techniques that external threats use involve network fingerprinting, scanning, exploitation attempts, and email phishing. Particular interest is given to phishing attacks by external threats because of their rate of success. Common phishing emails are sent to a large number of recipients. These emails instruct recipients to download attached files or navigate to a malicious link. They appear legitimate by using anticipated branding and letterhead.

By capturing historical data regarding certain groups within certain threat categories and by using techniques such as game theory, context can be created by which an organization can determine the likelihood of its assets being targeted and how they might be targeted. This situational awareness can translate to improved enterprise risk management decisions leading to improvements in areas such as organizational training and awareness, technology purchases and policy development.

A.3 Internal Threats

Insider threats have received the major amount of recent attention. Internal threat actors have the potential for disclosing not only classified or sensitive data to the organization but regulated privacy data about users or customers. Insider threats have access to internal information systems, security practices, knowledge of the organization vulnerabilities, and intellectual property.

Insider threats are often ex-employees or disgruntled current employees who believe the organization has done something wrong. Often, they see their actions as noble and appropriate for a given problem.

In response to insider threats the president, National Security Council and other organizational offices within the federal government have drafted and released guidance for how their own organizations should monitor and protect their systems against insider threats. These directives enumerate requirements for monitoring behavior and identifying potential threat actors. Common directives are listed below:

- Executive Order 13587, Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing and Safeguarding of Classified Information [White House 2011]
- Committee on National Security Systems Directive 504, Directive on Protecting NSS from Insider Threat [CNSS 2014]
- Intelligence Community Standard 700-2, Use of Audit Data for Insider Threat Detection [DoD 2006]

The SEI/CERT Insider Threat database identifies over 700 cases of malicious events involving insiders, dating back to the 1950s. Regardless of the insider threat's motivations, there are two primary effects achieved by their efforts. Sabotage and theft are the most common desired effects. By assessing the attack patterns of insider threats, better risk management decisions can be made leading to improvements in areas such as human-resource management, policy, and leader development.

A.4 Physical Threats

Physical systems must also be continually monitored to ensure threats don't affect security and availability. Often, physical threats are overlooked due to limited infrastructure and planning resources. The effects caused by physical threats are often more apparent than others.

Physical threats can be categorized in three ways:

1. hardware based
2. environmental
3. sustainability

Hardware-based threats include physical damage to servers, routing, and other components of the network. Electrical threats are often included with other hardware-based threats. These include availability and power line conditioning threats.

Environmental threats are related to the physical infrastructure of server and routing components. Temperature and humidity considerations must be designed into a facilities construction plan. Periodic reviews and facility validations provide mechanisms to ensure threat mitigation.

By capturing, analyzing, and reporting historical data on impacts caused by physical threats, improved risk management decisions can be made in areas such as facilities management, acquisition purchases, and policy development.

For all threat categories, risk analysts, by using techniques such as machine learning and text analysis, can extract certain threat categories' intentions and abilities for numerical comparison. By analyzing their previous and ongoing attack patterns, the analyst can use this measurement to provide a quantitative and qualitative aspect to the risk posture of an organization. If a particular threat, attack pattern, and target combination is more relevant to one organization than another, then this measurement can be used to heighten awareness for certain cybersecurity governance actions at that organization, such as training and awareness or a change in acquisition strategy.

Appendix B: Automated Vulnerability Collection

B.1 The Current State of Vulnerability Information Collection

Automated vulnerability management has evolved at a moderate pace over the last five years since the creation of the Security Content Automation Protocol (SCAP) standard. While many different processes and tools have been employed to support traditional vulnerability management, SCAP seeks to unify the industry's ability to identify and manage vulnerabilities in a common, automated fashion through the implementation of an open source identification format standard. SCAP has proven integral towards commonly identifying vulnerabilities (Common Vulnerability Enumeration – CVE), configuration setting mismanagement (Common Configuration Enumeration – CCE), and identification of assets (Common Platform Enumeration – CPE). SCAP is limited as a standard for formatting the identification of vulnerability data and does not address the approach for identifying vulnerabilities or organizations communicating the presence of those vulnerabilities. Additionally, the CVEs, CPEs, and CCEs are not all encompassing, as vendors have their own proprietary databases that evolve alongside the National Vulnerability Database (NVD). While this evolution has dramatically increased the ability to identify, manage, and mitigate vulnerabilities, a significant gap still exists in the ability to identify vulnerabilities elsewhere in the organization that do not pertain to software weaknesses.

This, in part, is due to an unsustainable approach to vulnerability management. At best, we default to “severity” ratings that are output from a tool or simply default to a CVSS score (where it exists). The total impact of the vulnerability (and the subsequent mitigation decisions) should instead be derived as a function of the

- Threat (likelihood of exploitation). This is calculated as a total of the following Exploitability Metrics, as a part of a CVSS score:
 - Access Vector (necessary threat actor proximity to vulnerability)
 - Access Complexity (non-authentication steps for exploitation)
 - Authentication (quantity of authentication steps)
 - Impact (based upon Confidentiality, Integrity, and Availability impacts, while taking into account the value of the asset relative to the D/A—the function or mission objective it performs)
- Vulnerability (identified weakness). This is the CVSS score calculated from the following variables:
 - Exploitability metrics
 - Impact metrics
 - Temporal metrics (ability to exploit and remediate, and confidence in understanding the vulnerability)
- Environmental metrics (collateral damage potential, percent of affected systems, and more focused confidentiality, integrity, and availability scoring)

In a traditional program, organizations might default to vulnerability only and not consider threat and impact when making mitigation decisions. Organizations should seek to create holistic, automated vulnerability detection and mitigation.

B.2 Recommended Way Ahead for Vulnerability Information Collection

Automating these capabilities is somewhat simplified within an environment that includes physical and logical privilege management, personnel security, credential management, and network architecture (as examples). Many organizations typically employ some policy for how they envision the aforementioned capabilities to work in their environment to support their mission or business objectives while maintaining an acceptable security posture, such as

- the specific privileges an individual needs to perform his or her job and the list of physical locations each individual can access
- the security clearance or suitability required for each position/role in an organization
- the credential or complexity required to access objects in an organization
- the list of known-bad software or other rules present on an access control device in the organization's network architecture

Each of the aforementioned examples is typically logically instantiated for enforcement in the production environment of an organization and is therefore measurable. Any deviations between the desired and actual state of instantiation of that policy should be identified as a vulnerability that can be measured. Due to the unique organizational technology and processes, as well as the corresponding vulnerabilities and respective impacts, the scoring of the subsequent risk must be qualitative as well as quantitative in order to impact decision making.

For instance, measuring the coverage of D/A systems by a data-loss-prevention capability can provide the cybersecurity governance decision maker with invaluable information regarding the organization's vulnerability to insider-threat actors. Another potential metric is the comparison of non-personal, elevated-privilege domain accounts approved by the network or organization's change control board with those actually in a system's directory services. Research by the authors of the *CERT® Insider Threat Guide* also indicates that actively managing the expectations of employees to minimize unmet disgruntlement can also reduce the likelihood of insider threat exploits. Measuring the degree to which position descriptions are current, especially those of 2210-series federal employees, can provide valuable insight. By coordinating with civilian personnel providers, cybersecurity governance decision makers can measure the average length of time since these position descriptions have been updated to ensure that any gaps are dealt with.

B.3 Challenges to Vulnerability Information Collection

Just like SCAP took some time to gain adoption from the technology vendors before the industry could adequately deploy an automated software vulnerability management process, many of the traditional legacy tools were not conceived to support such collection of data to compare policy to actual deployed instantiations of a capability.

Additionally, a tremendous amount of organizational and system correlation of data and processes is required to inform automated, analytics-backed, and visualized vulnerability management that

is comparable to organizational policy. This coordination includes the collection of vulnerability data from within network environments, analytics results for continuous improvement, as well as vulnerability feeds from outside organizations. The data itself is a combination of network visibility and hardware asset management, software vulnerabilities, and user, admin, and system account and access-related vulnerabilities. The processes supporting all of this are developed and improved upon through a governance structure for the management of metric collection and overall metric maturity. Governance must support the alignment of enterprise vulnerability policy, process, and risk scoring (the desired state of vulnerability management) with the deployed vulnerability detection systems and sensors (the actual state of vulnerability management) utilizing CVEs, US-CERT incident reports, and other vulnerability data sources. Utilizing the correlated risk scoring of vulnerabilities and balanced with the risk appetite of an organization, a qualitative vulnerability score can be defined as part of the governance ability to translate technical vulnerabilities into impactful risk management messaging for management.

Finally, the maturity of vulnerability management depends upon a sound governance process that accounts for prioritized risk identification, not only for the automation and visualization systems supporting the metrics, but also for the incremental deployment of any new systems and processes. The risks should be prioritized based upon a continuously improved and updated scoring based upon vulnerabilities and impact levels.

Appendix C: Use of Indices and Quantitative Methods to Enhance Cybersecurity Governance Decision-Making Insights and Lessons from Data Science Literature

C.1 Introduction

The state of research summarized in this section represents an initial look at literature from the past five years directly tied to security indices from different domains, coupled with a sample of foundational references related to the analysis and visualization of data. Although not comprehensive by any means, this research summary will provide the necessary guidance for the development of a governance index and serve as the foundation for ongoing and more comprehensive research into topics that will necessarily inform a governance index.

This summary focuses on five specific research investigations into security-related indices, each from a different industry or domain. In no particular order, the five investigations include

- an energy security index researched by Georgia Tech [Brown 2011]
- a financial services secrecy index researched by the Tax Justice Network [Christensen 2010]
- a space security index researched by the Canadian government, the Ploughshares Fund, and The Simons Foundation [West 2009]
- performance measures for U.S. domestic counterterrorism researched by Rand [Jackson 2009]
- a nuclear materials security index researched by the Nuclear Threat Initiative [NTI Security Index 2012]

This summary is presented in the following sections aligned to the following themes:

- a basic description of the nature of each index
- goals and outcomes pursued across the indices
- the methodology of developing the set of indices
- the data issues noted across the indices
- the analytics employed across the indices
- the key lessons learned from the five index experiences
- the noted gaps and conclusions from reflection on the five indices and other noted foundational references

C.2 Overview of Indices

Although the indices included in this summary are quite diverse, they demonstrate a number of interesting similarities and differences. The energy security index is defined as a measure of “equitably providing available, affordable, reliable, efficient, environmentally benign, proactively governed and socially acceptable energy services to end-users” [Brown 2011, p. 4]. This definition came from results of a literature search on the energy security topic in which dominant issues

were deduced from an occurrence standpoint in the literature [Brown 2011, p. 4]. The original motivation for this index arose from a “lack of a worldwide consensus about the nature of the energy security problem and an increased need to identify alternative approaches to address diverse countries” [Brown 2011, p. 4].

In similar fashion, the financial secrecy index arose from a desire to add transparency to the worldwide financial services industry [Christensen 2010, p. 1]. Developed in 2008, the financial secrecy index ranks worldwide jurisdictions on the opaqueness of their financial markets; for example, it ranks the world’s tax havens [Christensen 2010, p. 1].

Another index with an international flavor is the space security index. This index demonstrates a policy- and goal-driven approach to creating an index [West 2009, p. 1]. Developed by Canada, the Ploughshares Fund, and The Simons Foundation, the space security index is the only annual, comprehensive measure on activities in outer space and their measurable impact on security [West 2009, p. 1]. It defines space security as “secure and sustainable access to and use of space, and freedom from space-based threats” [West 2009, p. 1].

Closer to home, the performance measures associated with U.S. domestic counterterrorism intelligence focus on the mission effectiveness of preventing terrorist attacks originating from within the U.S. [Jackson 2009, p. 179]. Lastly, the NTI security index is a unique, public baseline assessment of the status of security conditions of nuclear materials around the world and evaluates a broad range of publicly available indicators of practices and conditions [NTI Security Index 2012, p. 6].

C.3 Goals and Outcomes Pursued by the Indices

All five security-related indices defined goals and outcomes at the outset of the development of the index. However, the diversity of goals and outcomes selected and how they were selected may offer insight to the approach for a cybersecurity governance index, both in the short term and long term. The NTI security index identified a need to measure risk, track progress, and hold states accountable with regards to nuclear materials security [NTI Security Index 2012, p. 6]. The authors of the index wanted not just a viewing console but a foundation for ongoing improvement with inclusion of recommendations [NTI Security Index 2012, p. 6]. They wanted the ability to set priorities and provide a system of assurance, accountability, and action [NTI Security Index 2012, p. 20]. From the outset, they embraced four principles for the index:

1. a robust analytical framework
2. an open and inclusive process
3. an international perspective
4. actionable policy prescriptions [NTI Security Index 2012, p. 24]

Likewise, the RAND discussion of performance measures for U.S. domestic counterterrorism intelligence began with a need to inform U.S. policy [Jackson 2009, p. 181] and a subsequent desire to support decision makers in assessing how doing something differently might actually be better [Jackson 2009, p. 179]. The RAND authors decided measures of both performance and acceptability were needed, with a specific focus on performance measures associated with internal efficiency and the monitoring of processes for producing intelligence outcomes [Jackson 2009, p.

179]. Lastly, they recognized the need to understand both the attack surface and the ability for the organization to react effectively and quickly to the attack [Jackson 2009, p. 183-184].

The space security index began simply with the desire to add transparency to the topic of space security. However, the authors of the index additionally wanted to raise critical questions, offer a new vision of security in space, and facilitate dialog on challenges and potential responses. From a policy perspective, they also wanted the index to be sensitive, selective, and effective at an acceptable cost. As a result, the authors turned their attention to measuring the capabilities and efficiencies of the organization [West 2009, p. 1].

The financial secrecy index began with a desire to model secrecy and scale of financial services simultaneously, while remaining objective and politically independent [Christensen 2010, p. 1].

The authors of the energy security index reviewed 91 peer-reviewed journal articles on the topic of energy security and used the frequency results to develop weightings of four dimensions that they derived from the journal articles: 1) availability, 2) affordability, 3) energy and economic efficiency, and 4) environmental stewardship [Brown 2011, p. 5-6].

The reader should note that measuring and modeling risk can be challenging. Therefore, the reader should be aware of some relatively recent publications on risk management, specifically the investigation by Douglas Hubbard into the failure of traditional risk management [Hubbard 2009]. Douglas Hubbard offers a number of recommendations for how to think of and measure risk that will be explored within the governance index context as needed.

C.4 Methodology for Developing the Indices

For most of the five indices, the methodology appears quite similar from a standpoint of identifying a hierarchy of categories or themes with subordinate measures or indicators. Within the NTI security index, five categories were identified, which were then defined by a subordinate set of 18 indicators and a total of 51 sub-indicators and associated weights [NTI Security Index 2012, p. 23]. The authors of the NTI index decided on a dynamic modeling approach in which some countries would be scored against all five categories while others would only and appropriately be scored against three of the categories [NTI Security Index 2012, p. 9]. Notably, the authors also declared that all countries should participate in the update and evolution of the index [NTI Security Index 2012, p. 10].

The RAND authors' methodology behind the domestic counterterrorism performance measures began with a desire to enable a "what-if" capability via the index, in which different scenarios could be evaluated and, using a baseline of desired performance for each organization and function, enable comparisons within and between entities across time [Jackson 2009, p. 186]. They anticipated that the index would also need to support a drill-down capability to facilitate action [Jackson 2009, p. 182]. To get started, they decided to first systematically think about "what the organization is designed to do and then how the organization is trying to do things, before deciding on appropriate measures" [Jackson 2009, p. 180]. As a consequence, they wanted measures of processes directly linked to outcomes [Jackson 2009, p. 180]. They identified five intelligence functions [information collection, sharing, analysis, storage, and action] that collectively produce the outcome of preventing terrorism [Jackson 2009, p. 181].

Notional performance measures were identified for each function, beginning with “first principles” and then with measures of how the functions work together [Jackson 2009, p. 181]. The authors looked for common measures that could be used to predict both outcomes of interest so that there would be a reduced risk of unhealthy tradeoffs [Jackson 2009, p. 182]. They decided data should enable anticipation of outcomes in measurable ways [Jackson 2009, p. 186] and that they wanted to measure the degree to which the data is acted upon [Jackson 2009, p. 193]. They did not let the ease or difficulty of measurement dictate what measures were needed. Rather, they initially disregarded that aspect and just looked at what metrics would be needed to support the desired analytics and predictions [Jackson 2009, p. 182]. Some other minor observations they made regarding their methodology include their recognition that feedback loops may be necessary to the modeling of interactions between different functions [Jackson 2009, p. 184]. Additionally, they decided to measure both interim and overall outcomes [Jackson 2009, p. 196] and to use sample measures identified by function, such as a measure of the capability, authority, and willingness to act [Jackson 2009, p. 186].

The methodology implemented to define the space security index began with the notion that the index would

- serve as a tool to help define the space security problem
- further identify the stakeholders
- set the stage for worldwide involvement
- raise questions for policy makers to address [West 2009, p. 1].

Eight indicators were employed to measure the following three broad areas of security:

1. the operating environment
2. actors
3. activities in space [West 2009, p. 1].

The methodology employed within the financial secrecy index is summarized for a quick read [Christensen 2010] and then detailed in 90 pages for those who wish to dig into the complete methodology and analytics [Tax Justice Network 2013]. This index, focused on transparency, decided to only use publicly verifiable information [Christensen 2010, p. 1]. They combined two measures: one qualitative and one quantitative. The qualitative measure is called the Opacity Score, which measures how aggressively a jurisdiction pursues secrecy features and “features likely to attract illicit financial flows.” The opacity score is measured using 15 indicators grouped under three themes:

1. transparency of ownership information
2. transparency of corporate activity
3. engagement in international cooperation to end harmful practices

The quantitative measure is a score of each jurisdiction’s volume of off-shore financial services offerings measured by either cross-border financial services trade or surrogate values of holdings in foreign portfolio assets [Christensen 2010].

The methodology employed by the energy security index was also quite simple. The authors decided on four energy security dimensions:

1. availability
2. affordability
3. efficiency
4. environmental stewardship [Brown 2011, p. 6]

The authors then decided they wanted to identify measurable indicators for each dimension, using only comparative indicators to enable

- setting energy targets
- measuring performance over time
- identifying tradeoffs and areas needing improvement [Brown 2011, p. 4]

Although they had access to a recent handbook that detailed over 1000 distinct metrics related to energy security, they decided to keep the index simple and chose a total of 10 indicators to measure the four dimensions [Brown 2011, p. 7]. Ironically, each of their indicators are inversely related to energy security (e.g., as the indicator goes down, the security index improves) [Brown 2011, p. 6]. Lastly, the authors chose to develop the models behind the index using sample data rather than expending effort to model with population data [Brown 2011, p. 8].

Beyond the five index research publications, another notable research source for methodology considerations for developing and evolving a governance index is a very popular book, *How to Measure Anything* by Douglas Hubbard [Hubbard 2010]. This specific reference has been used in SEI measurement training since its adoption in 2010, and provides a long list of tips and guidance on how to measure intangibles, as well as how to deal with real-world challenges in modeling uncertain factors. Participants developing governance indices would be strongly encouraged to obtain and read this reference.

C.5 Data Issues Experienced by the Indices

Only two of the five indices discussed data issues associated with the development or use of the index. The NTI security index authors decided that when sensitive data was not directly available, surrogate measures thought to correlate with the sensitive measures would be used [NTI Security Index 2012, p. 22]. The authors of the RAND report on domestic counterterrorism performance measures encountered a number of data issues as follows:

- issues related to signal versus noise in the data, along with resulting levels of false positives and negatives [Jackson 2009, p. 190]
- issues with accuracy and currency of the data [Jackson 2009, p. 191]
- issues with the shelf life of data [Jackson 2009 p. 191]
- issues with sensitivity and specificity amidst noisy data, often requiring corroborating data [Jackson 2009, p. 194]
- issues with attempted gaming of the index, requiring sensitivity to how much sharing and dissemination occurs with the combined data, foundational model, and reported results [Jackson 2009, p. 199]

A final note on challenges with data involves the use of subjective data, such as expert judgment. Although a governance index is intended to only use objective, concrete data, it may be necessary, in some early evolutionary version of a governance index, some transition point with a current security measurement system, or some desired longer term version of the governance index, to use expert judgment in specific areas. Recognizing this contingency, the Hubbard book called *How to Measure Anything* includes recent research and discussion of the fallacy of expert judgment, but continues to show how expert judgment can be calibrated through proper training [Hubbard 2010]. As a result, governance-index evolution incorporating expert judgment should assess application of the Hubbard training principles and warnings about expert judgment error.

C.6 Analytics Employed by the Indices

Only two of the five index authors shared analytics beyond simple arithmetic and computing averages. The authors of the financial secrecy index recognized that simple arithmetic weighted averages are not sensitive to signals that they wanted to pick up on. As a result, they adopted use of an arithmetic squaring of each jurisdiction's aggregate opacity score that ensured that minor differences in opacity among jurisdictions would be clearly evident. They also normalized the resulting opacity scores to a range of 0 to 100. The Opacity score is then simply multiplied by the quantitative score to get the index [Christensen 2010, p. 1-2]. Digging deeper in a separate methodology document brought the discovery that the authors assigned composite scores to the 15 indicators, each made up of a set of questions that are either Yes/No or on a scale of 1-4. There are also assigned weights to each of the questions in forming the composite score for the indicator [Tax Justice Network 2013].

The authors of the energy security index performed a wealth of rich analytics to confidently arrive at a compelling index. They conducted statistical correlation analysis between the metrics to understand which ones were correlated with which others. They also normalized their metrics to ensure they were not measuring effects of factors that they were not interested in. They considered two questions to determine external validity. First, do the indicators correlate with the dimensions? Second, do countries have similarities and differences in trends that align with the four dimensions? Both of these are answered by different forms of correlation analysis [Brown 2011, p. 7]. As the authors correctly recognized that correlation does not imply causality, they expended greater effort, involving domain experts to aid in judging the existence of causation. The authors also recognized that modeling the index could become problematic without distinguishing the two types of variation in the data:

1. special cause variation
2. common cause variation

They also decided to use z score transformations on all of the indicators before evaluating changes in the indicators and conducting the correlation analysis. Using such a transformation before conducting correlation and regression analysis is argued by some to be a more robust approach to correlation and regression [Brown 2011, p. 7]. The authors then utilized a different tool from the statistical toolkit, called factor analysis, in an attempt to show the relationship of the indicators to the four dimensions [Brown 2011, p. 15]. Next, the authors utilized another tool from the statistical toolkit called cluster analysis, to use country scores to see if distinct groupings of similar performing countries might exist. The cluster analysis was therefore used to characterize the countries

against the framework of the four dimensions and satisfy the earlier mentioned concerns of external validity [Brown 2011, p. 16].

C.7 Key Lessons Learned

Two of the five indices included discussion of some key lessons learned not already mentioned in other parts of this paper. The authors of the domestic counterterrorism performance measures noted that improper definition and use of measures would cause counterproductive behavior [Jackson 2009, p. 180]. They also found that the discussion of work functions and performance measures truly enriched the understanding of the overall organization work design and processes [Jackson 2009, p. 181]. During the development of the index, they had to renew emphasis on a balanced set of measures to minimize suboptimal behavior associated with the index [Jackson 2009, p. 182]. Lastly, the authors recognized that they had to address both reality and perception at times in the modeling and data collection [Jackson 2009, p. 199].

C.8 Noted Gaps and Conclusions

At this stage of investigation, the literature concerning security-related indices has been noted as weak in these two areas:

- lack of discussion of the visualization opportunities within the security index domain
- lack of serious use of modern statistical and probabilistic methods, albeit noting that one of the five indices (energy security index) did use some statistical analysis

The visualization of data options for a governance index will be enriched further by first revisiting the heritage of research into data visualization documented in a historical form [Friendly 2005], as well as through three seminal research books by the renowned author, Edward R. Tufte, covering the display of data [Tufte 2001], visual explanations [Tufte 1997] and envisioning information [Tufte 1990].

C.9 Implementation of a Cybersecurity Governance-Related Index

In this section, we discuss in layman's terms how to use statistical methods to create greater situational awareness in using an index, when facing real-world challenges, such as incomplete but continuously improving data. The governance index is intended to provide governance data in a manner that is approachable and actionable by executive and senior-level management, by utilizing state-of-the-practice quantitative methods for synthesizing data. Although a governance index is intended to provide a statistical and visual snapshot of the governance posture, its greater importance is as a starting point for drill-down to more granular levels of metric detail. The next sections describe specific capabilities to be designed into an index during its planned evolution.

C.10 Manual and Automatic Data Measurement

An index will inherently leverage data that is both manually acquired as well as automatically accessed via other automated systems and databases. To accomplish the required statistical analysis, the data will need to be collocated with statistical computing resources including statistical, probabilistic, and simulation tools. Data should be flagged as to source, enabling appropriate downstream analytical activities.

C.11 Quantitative and Qualitative Data Measurement

To ensure a robust approach to the formulation and operation of the index, we anticipate that both quantitative and qualitative data may require analysis. Therefore, text analytic tools may be employed to accomplish needed text analytics. Such analytics will enable analysis of key words, phrases, co-located words and phrases, sentiment analysis, and other measures useful to assess both the content and character of the qualitative data. Qualitative analytics may also produce quantitative factors useful in other quantitative and predictive modeling.

C.12 Data Sampling

Leveraging the ability of sampling statistics to infer outcomes without needing to collect, measure, and analyze all possible data, sampling of data in support of the index calculation will be identified when such sampling offers needed agility, shortened time and effort to collect data, or reductions in data storage requirements. Sampling statistics using Analysis of Variance (ANOVA), hypothesis testing, and other forms of more advanced multivariate analysis, such as structural equation modeling, may be employed to produce outcomes that otherwise would require analysis of all possible data. Such sampling would involve power analysis, discussion of acceptable false positive and false negative rates, and statistical determination of minimum sample sizes.

C.13 Determining Data Quality, Integrity, and Completeness

Data quality and integrity checks may be accomplished in a number of ways statistically. Identification of outliers would be conducted, using techniques that may be automated and producing flags to data owners for investigation. Additionally, statistical techniques should be used to determine the likelihood that suspect data streams are “gamed” or otherwise artificially created using analytics derived from Benford’s Law, otherwise known as the First Digit Law. This law essentially has analyzed how often each digit appears in the first, second, and third location (and so on) of a number reported within different contexts. This law is also actively used by the U.S. Department of Defense to identify suspect reporting in data reports for program earned value management.

C.14 Determining Signal Versus Noise in Data

For most of the root level and composite data measures feeding the index, statistical analysis may be used to help determine whether the fluctuation of a given measure across time is “noise” (e.g., statistically expected variation) or a “true signal” (e.g., statistically unexpected variation). Use of this signal versus noise analytics will enable evaluation of a governance index and its subordinate measures for significant root cause changes that are affecting the index behavior that should be investigated (e.g., a change in a security analyst).

C.15 Determining Capability of Data Behavior to Comply With a Norm or Policy

Because most, if not all, data streams feeding the index will be evaluated as a distribution of behavior comprising sets of data points, we should also build upon the signal to noise analytics by using statistical process control chart techniques. These technique are used to determine process

“capability” or general usage of statistical confidence and prediction intervals for future data values, to determine if future behavior of data points will remain within arbitrary accepted “limits” specified by policy or internal guidance. The purpose of this focus is to identify in advance when one or more measures feeding the index may go outside of acceptable limits in future time periods. This would provide one way of enabling stakeholders to anticipate unacceptable performance of measures and/or the index in the future.

C.16 Investigating Time Trends and Cycles in Data

The many measures feeding the index may have periodic cycles across time for various reasons. We will be able to analyze and identify those measures and conduct time series analysis with possible time series auto correlation and regression, thereby capturing the behavior with regression equations. In this fashion, the cyclic data may be evaluated for signal to noise, capability, and predictive analytics using such regression equations. This evaluation will also prevent an index and its cyclic components from being misinterpreted by stakeholders.

C.17 Conducting Correlation Analysis Within Data

In disregarding whether subordinate data feeding the index is nominal, ordinal, interval, or ratio data, we will be able to conduct the appropriate statistical correlation analysis and provide stakeholders with correlation analysis, thereby informing them of the degree to which index component measures are correlated. Although not all conclusions of correlation should be viewed as a basis for concluding cause and effect, knowing that index components are correlated could be quite informative to stakeholders seeking more information on how to improve performance in the index and its components. Correlation results coupled with expert domain knowledge would provide a compelling case for action.

C.18 Comparing Organization Performance (Within and Between)

Statistical methods for making comparisons between agencies for a given index component measure will be accomplished using traditional statistical hypothesis testing. Hypothesis tests may be programmed to also compare one agency against itself across time. Different forms of analysis of variance (ANOVA) may be used to inform stakeholders about whether there are differences among organizations in context of one, two, or more index component measures. Such analysis most likely would be used to answer more sophisticated questions of organization comparisons, such as determining if organizations are different in context of two different controls.

C.19 Diagnosing Conditions Leading To Current Data Observations

Probabilistic modeling, such as Bayesian Belief Networks (BBNs), may be employed to help understand probabilistically what are the most likely conditions or related component behaviors that would lead to a given index or component behavior. This ability would support stakeholder root cause or proximate cause analysis of specific outcome levels that might be seen in the index or its subordinate components. Again, this analysis would support stakeholder research into corrective action planning.

C.20 Predicting Future Data Outcomes Using One or More Data Leading Indicators

To enable the index to not only look backwards to what has occurred or what is occurring in the index and its component measures but to also look forward and anticipate future behavior of the index and its component measures, we will build in predictive analytics using both statistical and probabilistic modeling. In this fashion, the index and associated scorecard would contain a collection of predicted performance values of the index and component measures for the enterprise and subordinate agencies. The intent with the predictive analytics would be to provide sufficient leading indication of index behavior that organizations could take pre-emptive measures to improve index performance.

C.21 Visualizing and Benchmarking Data Behavior and Performance

An index should comprise statistical data representing distributions rather than single data points. With that perspective, the index would make use of visual representations that depict distributional data including the mean, median, and percentile information. This representation will enable stakeholders to observe both upside and downside risk related to variation in the data. Additionally, boxplots and confidence intervals will be used so that stakeholders will not reach improper conclusions when viewing a single data point of performance.

C.22 Visual Requirements for the Statistical Data Associated with the Index

The index can be visually presented as a scorecard with selectable alternate views. A high-level heat map of an entire organization or enterprise will make up the highest level view. Views also can be selected down to a single organizational component and to a heat map view of the lowest level measures of the index for that component. Pivot table query capability would exist as well as views enabling “what-if” scenarios related to index component measures. Other views could include query mechanisms, akin to Siri, through which stakeholders ask questions and receive answers. Finally, standard reports could be requested as well as user-defined reports at the point of interface with the index scorecards.

C.23 Additional Use Consideration

These additional development principles for an index were identified through a literature search of other recent implementations of security indices:

- Ensure that the index computation does not enable specific weak performance to measures to be masked via the traditional approaches of averaging and weighted averaging.
- Statistically standardize the scoring of measures to safeguard the ability of the index to always differentiate performance among organizations, as well as to build in an “inflation” factor for performance, acknowledging that acceptable baselines of performance will increase over time.
- Accommodate the deletion and addition of measures within the index over time, especially in light of possible evolution measurement regimens and expected adoption curve by parts of the organization.

- Incorporate a stakeholder value system within the index through a framework of weights both at the individual measure and at the department and agency level, when performing roll-ups.

C.24 Analytical Solution for Such an Index

This use case will describe the generic use case and calculation of a governance index that may be implemented at any organizational level or for any selected grouping of departments and agencies, including hierarchal consolidations.

C.24.1 Preconditions

The following are conditions that must be met prior to successful implementation of a cybersecurity governance-related index:

- A selected group of organizational components is defined, including any hierarchical consolidations, for which the governance index shall be calculated.
- A selected set of data measures and sources is defined and made available for the governance index calculation.
- A point in time is defined for the governance index calculation.

C.24.2 Postconditions

The following conditions are expected to be met by the successful implementation of a cybersecurity governance-related index:

- A governance index value is computed for each selected organizational component.
- A governance index value is computed for each organizational consolidation level of the selected organizational components.

C.24.3 Value Statement

The audience for the requested governance index reporting will be able to observe the latest governance index values for selected organizational components and subsequent hierarchical consolidations. The governance index will reflect relative performance for a given set of selected governance index measures at a given point in time. This information may be used to assess governance index performance relative to other organizational components, as well as trending across time.

An individual desiring a governance index status report requests such a report by selecting a menu choice for governance index and answering system prompts. The system provides a prompt window asking for the following information: effective date of the governance index, identity of the organizational components and associated consolidations desired for the governance index, and identity of the categories and/or specific measures to be included in the governance index calculation. The system retrieves data associated with the set of identified measures as of the effective date specified for the identified organizational components. The system retrieves the associated default weights associated with the identified measures and corresponding consolidation levels of the selected organizational components.

Using default weights, the system recomputes all relative weights to ensure weights sum up to 100% for each level of consolidation, since not all eligible items at a given consolidation level

may have been selected for the requested governance index report. For each measure within a given organizational component, the system will compute the standard score (z score) of the measure for the organizational component as a function of a) the component value, b) the average value across all components, and c) standard deviation value of this measure across all organizational components within scope of this governance index report. For each computed standardized score (z score), the system will compute the corresponding standard score on the scale of 0-100%; any values of 0 will be subsequently and arbitrarily set to a value of 1% to enable use of the Weighted Product Model calculation. At this point, the system will begin computing consolidated standard scores of 1-100% beginning at the lowest level of consolidation and working up to the highest level of consolidation. The system will calculate each consolidation value using the Weighted Product Model, which may be seen in the following formula. The consolidation value, represented as $P(A_K)$, will be computed as follows, where n is the number of values to be consolidated and w is the weight assigned to each value:

$$P(A_K) = \prod_{j=1}^n (a_{Kj})^{w_j}, \text{ for } K = 1, 2, 3, \dots, m.$$

Figure 12: Weighted Product Model

As shown in the formula above, each value to participate in the consolidation (a_{kj}) will be raised to a power equal to the weight assigned to that value. In turn, this result is multiplied against the similar results of the other values to be consolidated to arrive at the overall consolidation value. The system will consolidate upward until the overall governance index for all the selected organizational components is calculated and then displays the overall top level governance index value to the user.

The user selects the overall governance index and asks for the consolidated governance index values one level down, then iteratively requests more drill down as desired. The system provides the consolidated governance index values one level down from the last level selected by the user. Each time the user selects a lower level index value, the system iterates by providing the consolidated index values one level down from the last level selected by the user. The user identifies any given measure or consolidation value within a given organizational component and asks for the time trend of that value. The system then reproduces the calculations of the entire request for all available historical months of data, based on the organizational components and selected measures and consolidations represented in the original request that is now part of the time trend request.

C.24.4 Implementation Considerations

The development of a governance index will be a journey rather than a destination. That is, it will be developed in iterations as new data sources become available or are refined, and as new governance index analytical capabilities are developed. A rigorous change-control process must be implemented to ensure that changes to the index and its display are managed in a methodical, pragmatic, yet agile fashion.

A series of capabilities ranging from simple trends to more advanced optimization and sensitivity capabilities will be mapped to an evolutionary plan for the governance index development. The breakout of the capabilities across time will now be clarified. Additional early work on the index includes the acquisition of text analytic tools to demonstrate the possibilities of feeding the index with measures computed by modern text analytics tools. Acquisition of additional modeling tools, including a tool to conduct structural equation modeling, is planned to support anticipated modeling of the index and its computation.

Additional literature search of detailed alternatives to the computation of the index should be accomplished and should result in a prototype simulator that demonstrates the combined use of the “Weighted Product Model” and a standardized scoring mechanism for individual index measures. Additional team brainstorming can confirm the approach of incorporating stakeholder value systems via a weighting system at both the individual measure level and at the organizational component level during index roll-ups. As work on the index proceeds, the team should ensure the index quantitative framework will be capable of adding and subtracting measures over time. This approach will support efforts to further streamline measures to be retained based on feedback from organizational components. It is envisioned that the following capabilities would be demonstrated first:

- trends
- signal versus noise
- capability against an established norm
- comparative analysis between organizational components

In a second phase, the governance index development should continue to evolve with the addition of the following capabilities:

- diagnosing the conditions leading to currently observed performance against one or more of the index measures
- predicting future performance of agencies and departments against one or more index measures and overall index performance
- characterization and correlation studies of index measures’ performance in support of exploratory analysis of agency and department performance

In a third phase, the maturation of the data collection and analysis framework should enable our conduct of empirical structural equation modeling, which would provide a solid, statistical basis for a more compelling quantitative approach to the index that would also identify measures that are insignificant and thus, candidates for elimination from future reporting for index purposes. In this way, the index would become more efficient in required measures.

In a fourth phase, the governance index development should continue to evolve with the addition of the following capabilities:

- monetization of existing organizational component performance as well as possible courses of improvement actions
- optimization analysis through improvements under a set of identified constraints within one or more organizational components

- complete “what-if” sensitivity analysis enabling management at the organizational component level and above to evaluate alternative courses of action to improve governance

In a fifth phase, full implementation of structural equation modeling should be possible with operationalization of the results in a Bayesian Belief Network (BBN) model. Such a BBN model will enable richer what-if analysis by capitalizing on the ability to accommodate missing or uncertain data, as well as modeling both historical data and expert subjective judgment.

C.25 Identifying New Sources of Data

To maintain the index and keep it current with the changing governance landscape, the sources of data for the index should be re-evaluated through a number of mechanisms:

- feedback from key organizational components pertaining to the index effectiveness
- comparison of fit of the organizational component index as compared to actual organizational component performance
- other key stakeholder feedback

As each potential new source of data is identified, an evaluation of the data quality, integrity, historical behavior, and degree of implementation should occur. This would be vital to properly integrating the new data sources into the index analytics.

C25.1 Integrating New Data Sources into the Index

The planned approach for the calculation of the index revolves around a hierarchy of data components and consolidations that involve a weighted, multiplicative scheme. This allows new data sources and components to be added in almost a plug-and-play mode. However, each time new data sources and components are added, the weighting scheme of components at each consolidation level may need to be revisited with stakeholders. Methods such as the Analytic Hierarchy Process (AHP) and fractional factorial orthogonal arrays may be used to elicit expert inputs enabling conversion to computed weights. Finally, the entire index calculation should be replicated within a simulation tool (discrete event or systems dynamics) to then run simulations of the component measures and observe/evaluate resulting index calculations and final organizational component index scores.

C.25.2 Re-Assessing the Data Sources Used for the Index

Over time, statistical analysis can offer efficiencies in the data sources and components needed for the index. Statistical methods exist to compare the informational value of each data source and component in context of the set of data sources and components in use. These same methods may be used to help prune the data sources and components to minimize the data information redundancy within the index calculation. Consequently, methods would enable a smaller, more efficient set of data to feed the index.⁹

C.25.3 Archiving Historical Index Component Data

Data related to organizational component index measure components and calculations can be archived for ongoing use in the maintenance and evolution of the index. By keeping a historical record of both the raw data feeding the index and the subsequent standardized scores, analysts are

able to conduct longitudinal studies that would further inform index calculations. This approach would be akin to the studies across time conducted to maintain the integrity, validity and reliability of standardized exams such as the American College Test (ACT) and the SAT.

C.25.4 Continuous Validation of the Index Results

Statistical methods can be used in a continuous or periodic fashion to validate the index calculations and organizational component scores when compared to actual organizational component performance. This validation activity will be most significant for the predictive analytics portion of the index. By periodically revisiting the set of statistical regression and probabilistic models within the index, shifts across time of performance, underlying influential factors and weight of index components may be recognized and accounted for. Domain experts associated with each data source will play a significant role in anticipating changes that would warrant revisiting each of the regression and probabilistic models.

Appendix D: Mapping of Facets of Cybersecurity Governance to Other Frameworks

Governance Facets	NIST CSF	CERT-RMM/CRR/C2M2	DOD JCIDS (DOTMLPF)
Doctrine and strategy	Inherent in all CSF tiers. ID.BE-2, ID.BE-3, ID.GV-1 are the most related subcategories.	Enterprise Focus (EF), specifically EF:SG1	Doctrine
Enterprise portfolio management	No Direct Mapping. ID.AM-4 is the most related subcategory.	Asset Definition and Management (ADM); Specifically ADM:SG1.SP3, ADM:SG2:SP1 and ADM:SG2:SP2	No Mapping
Financial resource management	No Mapping	Financial Resource Management (FRM) Enterprise Focus (EF), specifically EF:SG3.SP1	No Mapping
Enterprise acquisition and materiel management	No Direct Mapping. ID.BE-1 is the most related subcategory.	GG2.GP3 - Providing adequate resources for multiple process areas	Materiel
Human resources management and leader development	No Mapping	Human Resources Management (HRM), specifically HRM:SG2, HRM:SG3, and HRM:SG4 People Management (PM), specifically PM:SG3.SP1 and PM:SG3.SP2	Personnel Leadership
Organizational structure management	No Mapping	Human Resources Management (HRM), specifically HRM:SG1 Enterprise Focus, specifically EF:SG2.SP1 and EF:SG2.SP2 GG2.GP4 - Assigning process responsibility and authority for multiple process areas	Organization
Organizational training and awareness	Awareness and Training (PR.AT)	Organizational training and awareness (OTA)	Training
Legal, regulations, policy, orders, investigations, and compliance	No Direct Mapping. Most related categories/subcategories are ID.BE, specifically ID.BE-5. ID.GV, specifically ID.GV-1 and ID.GV-3	Resilience Requirement Development (RRD), specifically RRD:SG1.SP1 Resilience Requirement Management (RRM) Compliance (COMP)	No Mapping
Enterprise risk management	Inherent in all CSF tiers. Risk Assessment (ID.RA) and Risk Management Strategy (ID.RM)	Risk Management (RM)	No Mapping

Bibliography

[Allison 1999]

Allison, G. T., & Zelikow, P. *Essence of Decision: Explaining the Cuban Missile Crisis* 2nd ed. (Kindle Edition, location 3235). Longman, 1999.

[Angerman 2004]

Angerman, W. S. *Coming Full Circle with Boyd's OODA Loop Ideas: An Analysis of Innovation Diffusion and Evolution* (unpublished master's thesis). Air Force Institute of Technology, 2004.

[Boyd 1976]

Boyd, John. *Destruction and Creation*. 1976.
http://www.goalsys.com/books/documents/DESTRUCTION_AND_CREATION.pdf

[Boyd 1996]

Boyd, John R. *The Essence of Winning and Losing*. 1996. https://fasttransients.files.wordpress.com/2010/03/essence_of_winning_losing.pdf

[Boyd 2006]

John Boyd. *The Strategic Game of ? And ?* 2006.
http://www.dnipogo.org/boyd/strategic_game.pdf

[Brown 2011]

Brown, M. A., Sovacool, B. K., Wang, Y., & D'Agostino, A. L. *Energy Security Dimensions and Trends in Industrialized Countries*. 2011. <https://smartech.gatech.edu/handle/1853/41816>

[Capelli 2012]

Capelli, D., Moore, A., & Trzeciak, R. *The CERT Guide to Insider Threats* (Kindle Version). Addison-Wesley, 2012.

[Caralli 2010]

Caralli, R. A., Allen, J. H., & White, D. W. *CERT® Resilience Management Model: A Maturity Model for Managing Operational Resilience* (CMU/SEI-2010-TR-012). Software Engineering Institute, Carnegie Mellon University, 2010. <http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=9479>

[Caralli 2012]

Caralli, R., Knight, M., & Montgomery, A. *CERT Maturity Models 101: A Primer for Applying Maturity Models to Smart Grid Security, Resilience, and Interoperability*. Software Engineering Institute, Carnegie Mellon University. 2012. <http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=58916>

[Christensen 2010]

Christensen, J. "Measuring Secrecy." *The American Interest* 5, 6 (July-August 2010): 66-67.
<http://www.the-american-interest.com/2010/07/01/measuring-secrecy/>

[CNSS 2014]

Committee on National Security Systems. *CNSSD No. 504 Directive on Protecting NSS from Insider Threat*. February 24, 2014

[Covey 2013]

Covey, S. R. *The 7 Habits of Highly Effective People* (Kindle Edition). Simon & Schuster, 2013.

[Deming 1986]

Deming, W. E. *Out of the Crisis* (Kindle Edition). Massachusetts Institute of Technology, Center for Advanced Engineering Study, 1986.

[DOD 2006]

Department of Defense. *DoD Instruction 5240.05, Technical Surveillance Countermeasures (TSCM) Program*. February 22, 2006.

[DOE 2014]

U.S. Department of Energy. *Cybersecurity Capability Maturity Model*. DOE, February 2014.
<http://energy.gov/oe/cybersecurity-capability-maturity-model-c2m2-program/cybersecurity-capability-maturity-model-c2m2>

[DHS 2011]

Department of Homeland Security. *Blueprint for a Secure Cyber Future*. 2011.
<http://www.dhs.gov/xlibrary/assets/nppd/blueprint-for-a-secure-cyber-future.pdf>

[Doctrine 2014]

Doctrine Definition Merriam-Webster (Accessed April 25, 2014).
<http://www.merriam-webster.com/dictionary/doctrine>

[Fan 2013]

Fan, J., Han, F., & Liu, H. *Challenges of Big Data Analysis*. 2013. arXiv preprint arXiv:1308.1479

[FedScope 2014]

FedScope. *IBM Cognos PowerPlay Studio—Employment Trend*.
<http://www.fedscope.opm.gov/ibmcognos/cgi-bin/cognosisapi.dll> (2014).

[Friendly 2005]

Friendly, M. "Milestones in the History of Data Visualization: A Case Study in Statistical History." In C. Weihs, & W. Gaul, *Classification — The Ubiquitous Challenge* (pp. 34-52). Springer Berlin Heidelberg, 2005.

[GAO 2010]

Department of Homeland Security. Quadrennial Homeland Security Review Report (2010). DHS Transcript for: Assessing the Nation's Cybersecurity Strategy (GAO, 2013) Downloaded April 25, 2014 from <http://www.gao.gov/assets/660/652210.txt>

[Glenny 2011]

DarkMarket Cyberthieves, Cybercops, and You. (Kindle Edition) Alfred A. Knopf, 2011.

[Glenny 2012]

Glenny, M. *Darkmarket: How Hackers Became the New Mafia* (Kindle Edition). Vintage Books, 2012.

[Hale 2012]

Hales, Brad. *Cybersecurity – A Practical Approach to Actionable Intelligence*. Solarwinds, 2012.

[Hubbard 2009]

Hubbard, D. W. *The Failure of Risk Management: Why it is Broken and How to Fix it*. John Wiley & Sons, 2009.

[Hubbard 2010]

Hubbard, D. W. *How to Measure Anything: Finding the Value of “Intangibles” in Business*, 2nd ed. John Wiley & Sons, 2010.

[Hypothesis 2014]

Hypothesis Definition Merriam-Webster (Accessed June 20, 2014). <http://www.merriam-webster.com/dictionary/hypothesis>

[Jackson 2009]

Jackson, B. A. “Exploring Measures of Effectiveness for Domestic Intelligence: Addressing Questions of Capability and Acceptability.” In B. A. Jackson (Ed.), *The Challenge Of Domestic Intelligence in a Free Society* (pp. 179-204). 2009.
http://www.rand.org/content/dam/rand/pubs/monographs/2009/RAND_MG804.pdf

[Mandiant 2011]

Mandiant. APT1. *Exposing One of China’s Cyber Espionage Units*. 2011. http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf

[Military Personnel 2014]

Military Personnel—Reports Related to Military Personnel Statistics.
<https://www.dmdc.osd.mil/appj/dwp/reports.do?category=reports&subCat=milActDutReg> (2014).

[Nielsen 1990]

Nielsen, J. & Molich, R. “Heuristic evaluation of user interfaces,” 249-256. *Proceedings of the ACM CHI ’90 Conference*. Seattle, WA, April 1990. ACM, 1990.

[NIST 2010]

NIST SP 800-37. *Guide for Applying the Risk Management Framework to Federal Information Systems: A Life Cycle Approach*. Department of Commerce (NIST), 2010.

[NIST 2011]

NIST SP 800-39. *Managing Information Security Risk: Organization, Mission, and Information System View*. Department of Commerce (NIST), 2011.

[NIST 2012]

National Institute of Standards and Technology. *Framework for Improving Critical Infrastructure Security*. NIST, February 2014. <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214-final.pdf>

[NTI 2012]

NTI Nuclear Materials Security Index. *Building a Framework for Assurance, Accountability, and Action*. 2012. <http://ntiindex.org/wp-content/uploads/2014/01/2014-NTI-Index-Report.pdf>

[Olson 2012]

Olson, P. *We are Anonymous: inside the hacker world of Lulzsec, Anonymous, and the global cyber insurgency* (Kindle Version). Little, Brown and Co., 2012.

[Osinga 2007]

Osinga, Franz P. B. *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, 2007.

[Rustad 2009]

Rustad, M. *Internet Law in a Nutshell*. West Academic Publishing, 2009.

[Tague 2004]

Tague, Nancy R. *The Quality Toolbox*, Second Edition, pp 219-223. ASQ Quality Press, 2004. Excerpted at ASQ, “Decision Matrix. <http://asq.org/learn-about-quality/decision-making-tools/overview/decision-matrix.html> (2014).

[Tax Justice Network 2013]

Tax Justice Network. *Financial Secrecy Index 2013 Methodology*. 2013. <http://www.financialsecrecyindex.com/PDF/FSI-Methodology.pdf>

[Triantaphyllou 1989]

Triantaphyllou, E. & Mann, S. H. “An Examination of the Effectiveness of Multi-Dimensional Decision-Making Methods: A Decision-Making Paradox.” *International Journal of Decision Support Systems* 5, 3 (Sep 1989): 303–312.

[Triantaphyllou 2000]

Triantaphyllou, E. *Multi-Criteria Decision Making: A Comparative Study*. Kluwer Academic Publishers (now Springer), 2000. ISBN 0-7923-6607-7.

[Tufte 1990]

Tufte, E. R. *Envisioning Information*. Graphics Press, 1990.

[Tufte 1997]

Tufte, E. R. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Graphics Press, 1997.

[Tufte 2001]

Tufte, E. R. *The Visual Display of Quantitative Information*. Graphics Press, 2001.

[U.S. Congress 2002]

United States 107th Congress (§ 3543, H.R. 2458). Federal Information Security Management Act. 2002. <https://www.congress.gov/bill/107th-congress/house-bill/2458>

[U.S. Navy 2012]

Secretary of the United States Navy. *Manual for the Operation of the Joint Capabilities Integration and Development System*. 2012.

http://jtc.fhu.disa.mil/jtc_dri/pdfs/jcids_manual_19jan12.pdf

[Walton 1988]

Walton, M. & Deming, W. E. *The Deming Management Method*. Perigee Books, 1988.

[West 2009]

West, J. “Reaching Out: New Approaches to Security in Space.” *Ploughshares Monitor* 30, 1 (Spring 2009): 6-8. http://ploughshares.ca/pl_publications/reaching-out-new-approaches-to-security-in-space/

[Wood 2012]

Wood, Greg. “Solving Business Problems with Dogfights and OODA Loops.” *Left of the Deadline*. May 24, 2012. <http://blogs.sas.com/content/anz/2012/05/24/what-dogfights-and-ooda-loops-can-teach-us-about-solving-business-problems/>

[White House 2011]

Executive Order 13587—Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing and Safeguarding of Classified Information, October 7, 2011.

<https://www.whitehouse.gov/the-press-office/2011/10/07/executive-order-13587-structural-reforms-improve-security-classified-net>

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE September 2015		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Improving Federal Cybersecurity Governance Through Data-Driven Decision Making and Execution			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Douglas Gray, Julia Allen, Constantine Cois, Anne Connell, Erik Ebel, William Gulley, Michael Riley, Robert Stoddard, Marie Vaughan, Brian D. Wisniewski				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2015-TR-011	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFLCMC/PZE/Hanscom Enterprise Acquisition Division 20 Schilling Circle Building 1305 Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER n/a	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) Although efforts are underway through Information Security Continuous Monitoring initiatives to improve situational awareness and risk mitigation at the operational level, the federal government must make better enterprise-level cybersecurity decisions in the shortest time possible. This report outlines an approach called Data Driven Cybersecurity Governance Decision Making. This approach leverages the Observe, Orient, Decide, Act (OODA) loop used by the U.S. Department of Defense to enable decision makers at the strategic levels of government to best set the conditions for success at the point of execution. To best target the unique considerations of enterprise decision makers, this report discusses the difference between cybersecurity governance and cybersecurity operations. Within this context, it describes best practices in collecting and analyzing authoritative data present in the federal space to develop a level of situational awareness tailored to decision makers' needs in a cybersecurity governance scorecard. Cybersecurity governance decision makers can leverage this enhanced situational awareness to support a data-driven decision-making process that targets root causes of the problems facing the Federal government enterprise. Finally, the report discusses key considerations to ensure success at the point of execution based on work performed in the Observe, Orient, and Decide phases of the OODA Loop.				
14. SUBJECT TERMS Cyber, Governance, OODA, FISMA, Decision Support, Federal, Executive			15. NUMBER OF PAGES 73	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89) Prescribed by ANSI Std. Z39-18
298-102